When Bitcoin Mining Pools Run Dry

A Game-Theoretic Analysis of the Long-Term Impact of Attacks Between Mining Pools

Aron Laszka[1], Benjamin Johnson[2], and Jens Grossklags[3]

[1] Institute for Software Integrated Systems, Vanderbilt University, Nashville, USA
[2] CyLab, Carnegie Mellon University, Pittsburgh, USA
[3] College of Information Sciences and Technology, Pennsylvania State University, University Park, USA

Abstract. Bitcoin has established itself as the most successful cryptocurrency with adoption seen in many commercial scenarios. While most stakeholders have jointly benefited from the growing importance of Bitcoin, conflicting interests continue to negatively impact the ecosystem. In particular, incentives to derive short-term profits from attacks on mining pools threaten the long-term viability of Bitcoin.
We develop a game-theoretic model that allows us to capture short-term as well as long-term impacts of attacks against mining pools. Using this model, we study the conditions under which the mining pools have no incentives to launch attacks against each other (i.e., *peaceful* equilibria), and the conditions under which one mining pool is marginalized by attacks (i.e., *one-sided attack* equilibria). Our results provide guidelines for ensuring that the Bitcoin ecosystem remains long-term viable and trustworthy.

Keywords: Bitcoin, Mining, Attacks, DDoS, Game Theory

1 Introduction

Conceived in 2008, Bitcoin is a cryptocurrency system which is controlled through an online communication protocol and facilitated in a decentralized fashion [15]. Bitcoin has experienced considerable growth in popularity and constitutes the dominant cryptocurrency [2]. It has also increasingly found adoption as a viable payment scheme in mainstream electronic commerce.

Despite setbacks, such as the closure of the Mt. Gox exchange, most stakeholders of the Bitcoin ecosystem have profited from its development and expect to benefit also in the future from the trust placed in the cryptocurrency. As such, participants in the Bitcoin ecosystem share a *common goal* with the improvement (or avoidance of the erosion) of trust of the currency system. Interactions in most economic systems usually involve such common but at the same time also *conflicting interests* [18]. In the Bitcoin ecosystem, such misaligned incentives are manifested in several ways.

Most centrally, the process of mining new bitcoins is organized in the form of a race in which the miner that solves a proof-of-work task first will be rewarded; all other miners will leave empty-handed. Mining involves a probabilistic element, so that not only the most powerful miner would win a particular round of this competition with certainty. Nevertheless, individual miners have found it beneficial to join forces in the form of *mining pools*. For example, averaging mining proceeds across many participants makes earnings more predictable. The specific setup of each mining pool typically differs across several dimensions which can be tangible (e.g., related to the computing and communication infrastructure) or intangible (such as reputation or details of the payout schemes). We term the sum of these factors the *attractiveness* of a mining pool.

However, the decentralized and quasi-anonymous nature of the Bitcoin ecosystem also lowers the bar for unfair competition in the form of different *attacks* which can benefit a malicious mining pool. First, attackers may abuse the resources of unsuspecting computer users for mining purposes through security compromises [8,16]. Second, attackers may attempt to redirect or siphon off mining capabilities from a competing pool [12]. Third, attackers may diminish the mining power of competing pools, for example, through Distributed Denial of Service (DDoS) attacks [20], or by exploiting specific weaknesses in the implementation of the procedure/software used by a particular pool.

The third dimension of unfair competition has been the focus of two recent research contributions. Vasek et al. provided empirical evidence that, during a two-year period, about 29% of all known mining pools had been the subject of at least one DDoS attack, and for mining pools above 5% share of hash rate, the likelihood of suffering an attack was 63%. They further characterized the types of pools that are more likely to be attacked [20]. Building on these findings, we developed a game-theoretic model in [9], which investigated the adversarial interaction between two representative mining pools that can choose between productive and destructive investments (i.e., computing power vs. DDoS attack on its competitor). We found that the relative size of the mining pools is a critical factor for the incentives to engage in attacks.

Both research studies primarily focus on the immediate impact of attacks on mining pools, i.e., the temporary shutdown of mining power and its payoff consequences. However, previous research on DDoS attacks in the context of electronic commerce has shown that the future (medium or long-term) impact of service unavailability can actually be more significant [6]. In the context of mining pools, individual members can permanently shift to an unaffected pool, which lowers the future prospects of the attacked pool. Even though the long-term impact of attacks can have a strong influence on the behavior of service providers, this phenomenon has not been studied from a theoretical perspective for DDoS attacks in general, or in the context of Bitcoin mining pools in particular [9].

In this paper, we develop a game-theoretic model that allows us to investigate the long-term impact of attacks against mining pools. Using this model, we study the conditions under which the mining pools have no incentives to launch attacks

against each other (i.e., *peaceful* equilibria), and the conditions under which one mining pool is marginalized by attacks (i.e., *one-sided attack* equilibria). Our results provide guidelines for ensuring that the Bitcoin ecosystem remains long-term viable and trustworthy.

The remainder of this paper is organized as follows. In Section 2, we discuss related work relevant to the context of DDoS attacks in networked systems. We describe our model and key assumptions in Section 3. We conduct our analysis and present numerical results/illustrations in Sections 4 and 5, respectively. We offer concluding remarks in Section 6.

2 Related Work

Decision-making in the context of security has been extensively studied using various game-theoretic approaches [10,14]. Of particular interest to our work are studies which address the incentives for adversarial behaviors. For example, Schechter and Smith [17] build upon the literature on the economics of crime to construct a model of attackers in the computer-security context. The authors derive penalties and probabilities of enforcement that will deter a utility-optimizing attacker, who evaluates the risks and rewards of committing an offense. Clark and Konrad propose a game-theoretic model with one defender and one attacker [4]. In their model, the defending player has to successfully protect multiple nodes, while the attacker needs to compromise only a single node. Fultz and Grossklags study the competition between multiple strategic attackers in different interdependent decision-making scenarios [5,7].

Previous economic work has improved our understanding of DDoS attacks and potential countermeasures. Focusing on defender behaviors, Christin et al. investigate the incentives of a group of bounded rational agents when they face the threat of being absorbed into a botnet, e.g., for the purpose of a DDoS attack [3]. In contrast, Liu et al. model attackers and work towards identifying DDoS attacker strategies in a specific case study [13]. Li et al. model the incentives of a botnet master to maintain a zombie network for the primary purpose of renting a sufficiently large subset to a DDoS attacker [11]. The authors investigate whether this business relationship can remain profitable if defenders can pollute the botnet with decoy machines (which lowers the effectiveness of a DDoS attack). In addition, there are several other research studies which are concerned with the organization of effective countermeasures against DDoS [19,21].

As discussed in the introduction, our current work draws on Vasek et al. who provided empirical evidence on the prevalence of DDoS attacks in the Bitcoin economy [20]. They showed, for example, that the size of mining pools is related to the probability of being targeted by an attack. Those findings motivated us to develop a game-theoretic model of attack behaviors between two mining pools which is, however, restricted to studying short-term effects of attacks [9].

3 Model

3.1 Overview

Our modeling framework is designed to capture two distinct effects of attacks against mining pools. The first, and most obvious is a short-term effect on the revenue of the attacked pool. While an attack is ongoing, the communication of the pool is disrupted, and hence the revenue decreases. The second effect is a longer term decrease in the size of the pool. Due to the myopic behavior of miners, an ongoing attack may cause some miners to permanently leave the attacked pool and mine for other pools.

In our previous paper [9], we focused on the first effect and did not take into account the second. In this paper, we extend our analysis to incorporate both effects using a sequential game played over an indefinite number of rounds. In each round, the choices of players result in both short-term revenue consequences, which affect player utilities, as well as long-term migration consequences, which affect the relative sizes of pools in the next round.

Miner migration is an interesting feature in itself. Our modeling framework assumes that there is some level of migration in each round, regardless of any attacks. That is to say there is a percentage of miners who are sufficiently fluid in their preferences that they re-evaluate their choice of pool each round. The remaining percentage of miners in a given round will continue mining for the same pool in the next round.

Our main focus in studying this model will be determining steady-state equilibrium strategies. These are strategies that stabilize the long-term migration effects, so that the players' sizes remain the same from round to round; and that also constitute best-response strategies for each player. Steady-state equilibria are consistent with what we observe in the Bitcoin ecosystem, where we observe little change in the relative sizes of pools from round to round.[4]

Table 1 summarizes the notations used in the model.

3.2 Players

Our game has exactly two players: a bigger mining pool B and a smaller mining pool S. Each pool has a base level of attractiveness which we parameterize with two constants A_B, A_S. We may interpret A_B (for example) as the percentage of fluid miners who will migrate to pool B in the next round.

In contrast to the attractiveness levels which are fixed for the duration of the game, each pool also has a current size for each round. For example, $s_B^{(k)}$ is the relative size of pool B in round k. Relative size is interpreted to mean the percentage of hash power possessed by its miners, compared to the entire Bitcoin ecosystem.

[4] Steady-state equilibrium analysis has been used relatively sparingly in the security economics literature (see, for example, [1]), while it is a frequently employed solution concept in other areas of economics.

Table 1: Table of Notations

Symbol(s)	Constraints	Description
M	$\in [0,1]$	Base miner migration rate
C	$\in [0,1]$	Unit cost of attack
A_B, A_S	$\in [0,1]$ and $A_B + A_S \in [0,1]$	Relative attractiveness of the pool
$s_B^{(k)}, s_S^{(k)}$	$\in [0,1]$ and $s_B^{(k)} + s_S^{(k)} \in [0,1]$	Relative size of the pool in round k
$a_B^{(k)}, a_S^{(k)}$	$\in [0,1]$	Attack level of the pool in round k

From round to round, the sizes of pools may change; and in fact it may happen that a bigger pool becomes a smaller pool in the next round. To be consistent with our terminology then, the salient feature we use to distinguish B from S is the assumption that $A_B \geq A_S$.

3.3 Choices

In each round, players simultaneously choose an attack level in $[0,1]$. In round k, pool B chooses $a_B^{(k)}$, while pool S chooses $a_S^{(k)}$. The attack level is intended to be interpreted generically, independent of the specific attack form. The attack could be distributed denial of service (DDoS), or any other form of adversarial action that disrupts the attacked pool's mining efforts.

3.4 Consequences

The choices of mining pools affect both the short-term utilities of players, as well as the longer-term size of each pool as a result of miner migration.

Short-Term Consequences Let $C \in [0,1]$ be the per unit cost of an attack, then the utility of pool B in round k can be expressed in terms of the relative sizes of B and S via

$$u_B^{(k)} = \frac{s_B^{(k)} \cdot (1 - a_S^{(k)})}{1 - s_B^{(k)} \cdot a_S^{(k)} - s_S^{(k)} \cdot a_B^{(k)}} - C \cdot a_B^{(k)} \ . \tag{1}$$

In the above formula, $s_B^{(k)} \cdot (1 - a_S^{(k)})$ is the mining power of B considering S's attack, $1 - s_B^{(k)} \cdot a_S^{(k)} - s_S^{(k)}$ is the mining power of the whole Bitcoin ecosystem considering both attacks, and $C \cdot a_B^{(k)}$ is the total cost of attack incurred by B.

The utility function is designed to correspond directly to the percentage of mining revenue obtain by the pool in the given round. The relative amount of coins being mined in round k is decreased by the two players' attacks, which explains the denominator; while the relative amount of coins being mined by

pool B in round k is affected proportionally to the the attack level against pool B by pool S, which explains the numerator.

Note that by symmetry, we have the utility of S in round k as

$$u_S^{(k)} = \frac{s_S^{(k)} \cdot (1 - a_B^{(k)})}{1 - s_S^{(k)} \cdot a_B^{(k)} - s_B^{(k)} \cdot a_S^{(k)}} - C \cdot a_S^{(k)} \ . \tag{2}$$

Long-Term Consequences Attack strategies also have long-term consequences, that do not affect the players' immediate revenue, but do affect miner migration, and hence the relative sizes of the pools in the next round.

In each round, the miners that are affected by an attack re-evaluate their choices and start to migrate. Formally, in each round, $s_B^{(k)} \cdot a_S^{(k)}$ (or $s_S^{(k)} \cdot a_B^{(k)}$) miners leave pool B (or S) due to attacks. The group of migrating miners redistribute themselves in the next round among the pools in a manner proportional to each pool's relative attractiveness level (A_B, A_S, or $1 - A_B - A_S$, for pool B, pool S, and all other pools, respectively).

Miners may also re-evaluate their choices from time to time even when they are not affected by an attack. Let $M \in [0, 1]$ denote the level of this base migration. If $M = 0$, then from any initial state $s_B^{(0)}, s_S^{(0)}$, the pool sizes remain constant from round to round when there are no attacks. If $M = 1$, then every miner re-evaluates her pool choice in each round. Similarly to migration due to attacks, the group of migrating miners redistribute themselves among the pools in a manner proportional to each pool's relative attractiveness level.

We may thus express the relative size of pool B in round $k + 1$ in terms of the relative sizes of pools in round k, together with attack levels and the base migration rate:

$$\begin{aligned} s_B^{(k+1)} = \ & s_B^{(k)} \\ & + A_B \cdot [(1 - s_B^{(k)}) \cdot M + s_S^{(k)} a_B^{(k)} (1 - M)] \quad \text{(migration into } B\text{)} \\ & - s_B^{(k)} \cdot (1 - A_B) \cdot [M + a_S^{(k)} (1 - M)] \quad \text{(migration out of } B\text{).} \end{aligned} \tag{3}$$

Analogously, the relative size of pool S in round $k + 1$ may be expressed as

$$\begin{aligned} s_S^{(k+1)} = \ & s_S^{(k)} \\ & + A_S \cdot [(1 - s_S^{(k)}) \cdot M + s_B^{(k)} a_S^{(k)} (1 - M)] \quad \text{(migration into } S\text{)} \\ & - s_S^{(k)} \cdot (1 - A_S) \cdot [M + a_B^{(k)} (1 - M)] \quad \text{(migration out of } S\text{).} \end{aligned} \tag{4}$$

4 Model Analysis

We begin the analysis of our modeling framework by proving a uniqueness result for each pool's relative size in a steady-state equilibrium.

4.1 Steady-State Pool Sizes

Theorem 1. *If $M > 0$, then for any strategy profile (a_S, a_B), there exists a unique pair of relative sizes (s_S^*, s_B^*) such that*

$$(s_S^{(k+1)}, s_B^{(k+1)}) = (s_S^{(k)}, s_B^{(k)}) = (s_S^*, s_B^*).$$

Proof. Given a strategy profile (a_S, a_B), the conditions of the theorem require (s_S^*, s_B^*) to satisfy

$$\begin{aligned} & s_B^* \cdot (1 - A_B) \cdot [M + a_S(1 - M)] && \text{(migration out of } B) \\ & = A_B \cdot [(1 - s_B^*) \cdot M + s_S^* a_B (1 - M)] && \text{(migration into } B) \end{aligned}$$

and

$$\begin{aligned} & s_S^* \cdot (1 - A_S) \cdot [M + a_B(1 - M)] && \text{(migration out of } S) \\ & = A_S \cdot [(1 - s_S^*) \cdot M + s_B^* a_S (1 - M)] && \text{(migration into } S). \end{aligned}$$

Solving B's migration equation for s_B^*, we obtain

$$s_B^* = \frac{A_B \cdot [M + s_S^* a_B (1 - M)]}{M + a_S(1 - A_B)(1 - M)}.$$

This linear constraint is satisfied by all points (s_S, s_B) on the line segment connecting the points $\left(0, \frac{A_B M}{M + a_S(1-A_B)(1-M)}\right)$ and $\left(1, \frac{A_B[M + a_B(1-M)]}{M + a_S(1-A_B)(1-M)}\right)$.

Similarly, the migration equation for S may be reduced to the linear constraint

$$s_B^* = \frac{-A_S M + s_S^*[M + a_B(1 - A_S)(1 - M)]}{A_S a_S (1 - M)},$$

which is satisfied by all pairs (s_S, s_B) on the line segment connecting the points $\left(0, \frac{-M}{a_S(1-M)}\right)$ and $\left(1, \frac{(1-A_S)[M + a_B(1-M)]}{a_S A_S (1-M)}\right)$.

We now wish to show that these two segments must intersect in the interval $[0, 1]$. More precisely, we claim that the second segment starts below the first segment when $s_S = 0$; and ends above the first segment when $s_S = 1$.

The two relevant inequalities are:

$$\frac{-M}{a_S(1 - M)} < \frac{A_B M}{M + a_S(1 - A_B)(1 - M)} \tag{5}$$

and

$$\frac{A_B[M + a_B(1 - M)]}{M + a_S(1 - A_B)(1 - M)} < \frac{(1 - A_S)[M + a_B(1 - M)]}{a_S A_S (1 - M)} \tag{6}$$

Inequality (5) follows because the first term is negative while the second term is positive (or zero if $A_B = 0$). Inequality (6) follows from:

$$\frac{A_B[M + a_B(1-M)]}{M + a_S(1-A_B)(1-M)} \leq \frac{(1-A_S)[M + a_B(1-M)]}{M + a_S A_S(1-M)} \quad (A_S + A_B \leq 1)$$
$$< \frac{(1-A_S)[M + a_B(1-M)]}{a_S A_S(1-M)} \quad (M > 0).$$

□

As a result of the theorem, the unique steady-state solution can be expressed directly in terms of the strategies a_S and a_B, the attractiveness levels A_S and A_B, and the migration constant M:

$$s_S^* = \frac{A_S M[M + a_S(1-M)]}{[M + a_B(1-A_S)(1-M)][M + a_S(1-A_B)(1-M)] - A_B a_B A_S a_S(1-M)^2} \tag{7}$$

$$s_B^* = \frac{A_B M[M + a_B(1-M)]}{[M + a_B(1-A_S)(1-M)][M + a_S(1-A_B)(1-M)] - A_B a_B A_S a_S(1-M)^2}. \tag{8}$$

Furthermore, we know that these values are in $[0, 1]$ under our modeling assumptions without having to do further case analysis.

4.2 Steady-State Pool Utilities

Since there is a unique pair of steady-state pool sizes for each strategy profile, we can find a steady-state equilibrium by assuming that the pool sizes and, hence, the players' utilities are given by Equations (7) and (8). In other words, given a strategy profile (a_S, a_B), we may write the utility of each pool under the assumption that the relative sizes are the steady-state sizes s_S^* and s_B^*.

For pool S, we obtain

$$u_S = \frac{A_S M[M + a_S(1-M)](1-a_B)}{Denominator} - a_S C, \tag{9}$$

and for pool B, we have

$$u_B = \frac{A_B M[M + a_B(1-M)](1-a_S)}{Denominator} - a_B C, \tag{10}$$

where

$$\begin{aligned} Denominator = \; & M + a_S(1-M-A_B)][M + a_B(1-M-A_S)] \\ & + a_S a_B[(1-A_S-A_B)(1-M)^2 - A_S A_B]. \end{aligned} \tag{11}$$

This formulation will permit us to determine all the steady-state equilibria for the sequential game presented in Section 3 by finding Nash equilibria for a single-shot two-player game with the above utilities.

4.3 Peaceful Equilibria

We begin our analysis of equilibria by determining the conditions under which it is a stable strategy profile for each player to refrain from attacking the other player.

Theorem 2. *The strategy profile* $(a_S, a_B) = (0, 0)$ *is a Nash equilibrium just in case*

$$C \geq \frac{A_B A_S}{\min\{M, 1 - A_B, 1 - A_S\}}. \tag{12}$$

Proof. Suppose that $a_S = 0$. We want to characterize the conditions under which $a_B = 0$ is a best response. First, we express the utility of B in a steady state by substituting $a_S = 0$ into Equation(10), obtaining

$$u_B = \frac{A_B[M + a_B(1 - M)]}{[M + a_B(1 - M - A_S)]} - a_B C \ . \tag{13}$$

When B does not attack ($a_B = 0$), her resulting steady-state utility is $u_B = A_B$; while if she attacks with full force ($a_B = 1$) her utility becomes $u_B = \frac{A_B}{1 - A_S} - C$. Because the utility function is analytic in a_B, any intermediate attack level can only be a utility-maximizing response strategy if the partial derivative of u_B with respect to a_B evaluated at that specific attack level is zero.

Computing the first and second partial derivatives of u_B with respect to a_B, we obtain

$$\frac{\partial u_B}{\partial a_B} = \frac{A_B A_S M}{[M + a_B(1 - M - A_S)]^2} - C \ , \tag{14}$$

and

$$\frac{\partial^2 u_B}{\partial a_B^2} = \frac{-2A_B A_S M(1 - M - A_S)}{[M + a_B(1 - M - A_S)]^3} \ . \tag{15}$$

Since the denominator of Equation (15) is always positive, the second derivative itself is of constant sign; and this sign is negative if and only if $1 - M - A_S > 0$, or equivalently $M < 1 - A_S$. It is only in this case where the roots of the first derivative will give relevant maximizing solutions for the attack level.

In the case $M > 1 - A_S$, the roots of the first derivative will give minimizing attack levels, and in the case $M = 1 - A_S$ the first derivative will be constant, and hence the utility B will be maximized at either one of the endpoints of $[0, 1]$, or on the entire interval.

Setting the derivative from Equation (14) equal to zero and solving for a_B, we obtain

$$a_B = \frac{\sqrt{\frac{A_B A_S M}{C}} - M}{1 - A_S - M} \ . \tag{16}$$

Now we consider two parameter cases.

– First, if $M \geq 1 - A_S$, then the maximizing attack level is one (or both) of the two endpoints 0 or 1. In this case an optimal response is $a_B = 0$ exactly when $A_B \geq \frac{A_B}{1-A_S} - C$, or equivalently, when

$$C \geq \frac{A_B A_S}{1 - A_S} = \frac{A_B A_S}{\min\{M, 1 - A_S\}}.$$

– Second, if $M < 1 - A_S$, then the maximizing attack level will be zero if and only if the point at which the derivative is zero is non-positive. Since we are in the case $1 - A_S < M$, we may deduce from Equation (16) that the analytically-maximizing a_B is non-positive if and only if $\sqrt{\frac{A_B A_S M}{C}} - M \geq 0$. This condition reduces to

$$C \geq \frac{A_B A_S}{M} = \frac{A_B A_S}{\min\{M, 1 - A_S\}}.$$

These two parameter cases exhaust all options; and we have shown that in each case $a_B = 0$ is a best response to $a_S = 0$ if and only if

$$C \geq \frac{A_B A_S}{\min\{M, 1 - A_S\}}.$$

To have $(a_S, a_B) = (0, 0)$ be an equilibrium, we also need $a_S = 0$ to be a best response to $a_B = 0$. By symmetry, this will happen if and only if

$$C \geq \frac{A_B A_S}{\min\{M, 1 - A_B\}}.$$

We conclude that a peaceful equilibrium exists if and only if

$$C \geq \frac{A_B A_S}{\min\{M, 1 - A_B, 1 - A_S\}}.$$

□

4.4 One-Sided Attack Equilibria

Our next special case to consider is when exactly one of the players attacks while the other remains peaceful.

Theorem 3. *The strategy profile $(a_S, a_B) = (0, 1)$ forms a Nash equilibrium if and only if*

$$C \leq \frac{A_B A_S}{(1 - A_S)^2} \cdot \min\{M, 1 - A_S\} \tag{17}$$

Proof. First suppose that $a_B = 1$. Then the utility of S is given by

$$u_S = -a_S C, \tag{18}$$

and this quantity is clearly maximized for $a_S \in [0,1]$ by taking $a_S = 0$. So $a_S = 0$ is always a best response to $a_B = 1$.

Next suppose that $a_S = 0$. We want to characterize the conditions under which $a_B = 1$ is a best response. From the previous special case analysis, we have

$$u_B = \frac{A_B[M + a_B(1-M)]}{[M + a_B(1-M-A_S)]} - a_B C \ .$$

Exactly as in the previous analysis, when $a_B = 0$, we have $u_B = A_B$; and if $a_B = 1$, we get $u_B = \frac{A_B}{1-A_S} - C$. The same conditions applied to the derivative of u_B determine when the maximizing value of a_B is reached at a boundary point (0 or 1) or whether it may relate to where the derivative is zero at

$$a_B = \frac{\sqrt{\frac{A_B A_S M}{C}} - M}{1 - A_S - M} \ .$$

The derivative value is relevant if and only if $M < 1 - A_S$.

In the case where $M \geq 1 - A_S$, a maximizing attack level is one of the endpoints of [0,1]; and in this parameter case, the best response is $a_B = 1$ if and only if

$$C \leq \frac{A_B A_S}{1 - A_S} = \frac{A_B A_S}{(1-A_S)^2} \cdot \min\{M, 1 - A_S\}.$$

In the case $M < 1 - A_S$, the maximizing value is 1 if and only if the global analytically-derived maximum is at least 1. This happens when

$$\begin{aligned}
\frac{\sqrt{\frac{A_B A_S M}{C}} - M}{1 - A_S - M} &\geq 1 \\
\sqrt{\frac{A_B A_S M}{C}} - M &\geq 1 - A_S - M \\
\frac{A_B A_S M}{C} &\geq (1 - A_S)^2 \\
\frac{A_B A_S M}{(1-A_S)^2} &\geq C \\
C &\leq \frac{A_B A_S}{(1-A_S)^2} \cdot \min\{M, 1 - A_S\}
\end{aligned}$$

Again these two parameter cases exhaust all options; and in each case $(a_S, a_B) = (0,1)$ is an equilibrium configuration if and only if

$$C \leq \frac{A_B A_S}{(1-A_S)^2} \cdot \min\{M, 1 - A_S\}.$$

□

To have $(a_S, a_B) = (1, 0)$ be an equilibrium, we would need the condition

$$C \leq \frac{A_B A_S}{(1 - A_B)^2} \cdot \min\{M, 1 - A_B\}.$$

Of course both conditions may be simultaneously satisfied for C sufficiently small, in which case both one-sided attack configurations will be equilibria.

5 Numerical Illustrations

5.1 The Peaceful Equilibrium

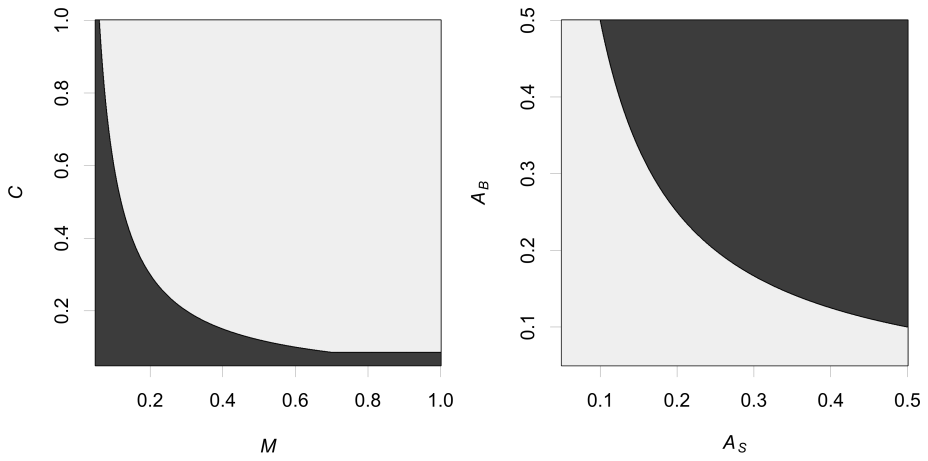

(a) $A_S = 0.2$, $A_B = 0.3$ (b) $M = 0.5$, $C = 0.1$

Fig. 1: Existence of the peaceful equilibrium (i.e., $a_S = 0$ and $a_B = 0$). Lighter shaded areas represent parameter combinations where the peaceful equilibrium exists.

Figure 1 shows the parameter combinations where the peaceful equilibrium exists (i.e., $a_S = 0$ and $a_B = 0$ being best responses to each other). In Figure 1(a), we fix the parameters $A_S = 0.2$ and $A_B = 0.3$, and we vary the parameters M and C. The figure shows that, if both M and C are high, then the peaceful equilibrium is possible; however, if either of these parameters is low, then there can be no peace. The latter is especially important in the case of M, which has no effect on the existence of the peaceful equilibrium once its value reaches $1 - A_j$ (0.7 in the figure). In practice, this means that both the pools and the users have to act in order to reach the peaceful equilibrium: the pools have to employ defensive countermeasures which increase C, while the users have to be proactive in their mining-pool choice, which increases M.

In Figure1(b), we fix the parameters $M = 0.5$ and $C = 0.1$, and we vary the parameters A_S and A_B. The figure shows that the peaceful equilibrium exists if either one (or both) of the pools has a low attractiveness. In practice, this means that a healthy competition between the pools, in which both of them try to attract miners, is very beneficial: not only will it result in better deals for the miners, but it may also bring peace.

5.2 One-Sided Attack Equilibria

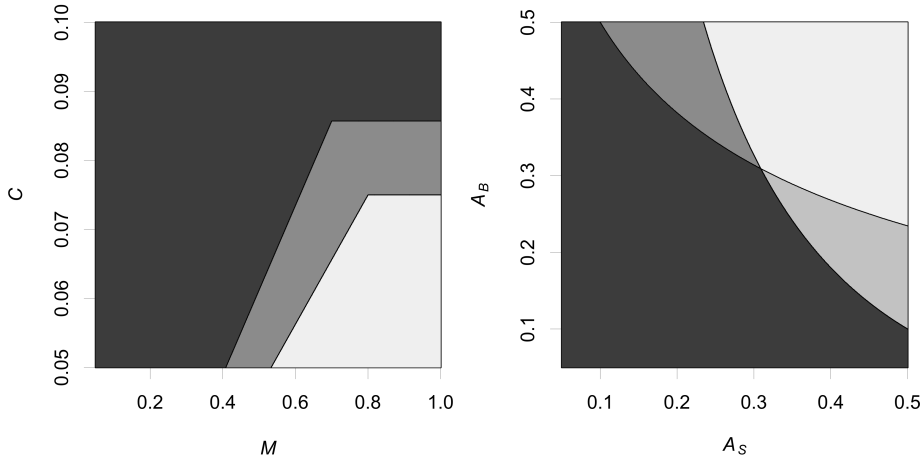

(a) $A_S = 0.2$, $A_B = 0.3$ (b) $M = 0.5$, $C = 0.1$

Fig. 2: Existence of one-sided attack equilibria. Shades from darker to lighter: no one-sided attack equilibrium, $(a_S = 1, a_B = 0)$ forms an equilibrium, $(a_S = 0, a_B = 1)$ forms an equilibrium, both form equilibria.

Figure 2 shows the parameter combinations where one-sided attack equilibria exist (i.e., when pool $i \in \{S, B\}$ playing $a_i = 1$ and the other pool playing $a_{\bar{i}} = 0$ forms an equilibrium). In Figure 2, we fix the parameters $A_S = 0.2$ and $A_B = 0.3$, and we vary the parameters M and C. The figure shows that one-sided attack equilibria are more likely to exist when M is high but C is low. This means that, to avoid a one-sided attack equilibrium, the pools must employ defensive countermeasures that increase the cost of attack C. Furthermore, the figure also shows that less attractive pools are more likely to play a one-sided attack strategy (darker shaded middle section representing $(a_S = 1, a_B = 0)$). While this may seem counterintuitive at first, it is actually very easy to explain. The more attractive pool has more miners even without launching an attack; hence, it is less inclined to dominate the other player with a marginalizing attack. The less attractive pool, on the other hand, has a lot to gain from such an attack; hence, it is more inclined to try to "steal" the miners of the more attractive pool.

In Figure 2(b), we fix the parameters $M = 0.5$ and $C = 0.1$, and we vary the parameters A_S and A_B. The figure shows that, once a pool is highly attractive to the miners, the other pool will have an incentive to launch a marginalizing attack against it. While attacks are generally harmful to the Bitcoin ecosystem, they have positive effects in this context, as they prevent one pool from growing too large.

6 Conclusion and Future Work

In this paper, we proposed a game-theoretic model of attacks between Bitcoin mining pools, which – to the best of our knowledge – is the first study to consider long-term consequences. The analysis of our model has revealed a number of interesting implications for making the Bitcoin ecosystem more viable. We have seen that, in order to make the peaceful equilibrium viable, the unit cost of attack and the miners' base migration rate both have to be increased, and no pool can have an overwhelming attractiveness. We have seen that these factors also help preventing a marginalizing attack against one pool.

We limited the mining pools' strategic choices to launching attacks and assumed that the effects of defensive countermeasures are incorporated into the unit cost of attack. In future work, we can extend this model by allowing the pools to deploy additional defenses for some fixed cost, which decrease the effectiveness of attacks. As another direction, we can also extend the model by allowing the pools to choose some of the parameters that affect their attractiveness levels. For example, a pool could choose to increase its fee, which decreases its attractiveness but increases its utility for a given steady-state size. Finally, in this paper, we modeled only two pools as strategic players, but we intend to extend our work towards the case of multiple mining pools.

Acknowledgments

We thank the reviewers for their detailed feedback. This work was supported in part by the National Science Foundation under Award CNS-1238959.

References

1. Bensoussan, A., Kantarcioglu, M., Hoe, S.: A game-theoretical approach for finding optimal strategies in a botnet defense model. In: Alpcan, T., Buttyan, L., Baras, J. (eds.) Decision and Game Theory for Security, Lecture Notes in Computer Science, vol. 6442, pp. 135–148. Springer Berlin Heidelberg (2010)
2. Böhme, R., Christin, N., Edelman, B., Moore, T.: Bitcoin. Journal of Economic Perspectives (forthcoming)
3. Christin, N., Grossklags, J., Chuang, J.: Near rationality and competitive equilibria in networked systems. In: Proceedings of the ACM SIGCOMM Workshop on Practice and Theory of Incentives in Networked Systems. pp. 213–219 (2004)

4. Clark, D., Konrad, K.: Asymmetric conflict: Weakest link against best shot. Journal of Conflict Resolution 51(3), 457–469 (Jun 2007)
5. Fultz, N., Grossklags, J.: Blue versus red: Towards a model of distributed security attacks. In: Proceedings of the 13th International Conference on Financial Cryptography and Data Security (FC). pp. 167–183 (Feb 2009)
6. Goldfarb, A.: The medium-term effects of unavailability. Quantitative Marketing and Economics 4(2), 143–171 (Jun 2006)
7. Grossklags, J., Christin, N., Chuang, J.: Secure or insure? A game-theoretic analysis of information security games. In: Proceedings of the 2008 World Wide Web Conference (WWW'08). pp. 209–218 (Apr 2008)
8. Huang, D.Y., Dharmdasani, H., Meiklejohn, S., Dave, V., Grier, C., McCoy, D., Savage, S., Weaver, N., Snoeren, A.C., Levchenko, K.: Botcoin: Monetizing stolen cycles. In: Proceedings of the 2014 Network and Distributed System Security Symposium (NDSS) (2014)
9. Johnson, B., Laszka, A., Grossklags, J., Vasek, M., Moore, T.: Game-theoretic analysis of DDoS attacks against Bitcoin mining pools. In: Böhme, R., Brenner, M., Moore, T., Smith, M. (eds.) Financial Cryptography and Data Security (Proceedings of the 1st Workshop on Bitcoin Research), pp. 72–86. Lecture Notes in Computer Science, Springer Berlin Heidelberg (2014)
10. Laszka, A., Felegyhazi, M., Buttyán, L.: A survey of interdependent information security games. ACM Computing Surveys 47(2), 23:1–23:38 (Aug 2014)
11. Li, Z., Liao, Q., Blaich, A., Striegel, A.: Fighting botnets with economic uncertainty. Security and Communication Networks 4(10), 1104–1113 (Oct 2011)
12. Litke, P., Stewart, J.: BGP hijacking for cryptocurrency profit (Aug 2014), available at: http://www.secureworks.com/cyber-threat-intelligence/threats/bgp-hijacking-for-cryptocurrency-profit/
13. Liu, P., Zang, W., Yu, M.: Incentive-based modeling and inference of attacker intent, objectives, and strategies. ACM Transactions on Information and System Security 8(1), 78–118 (Feb 2005)
14. Manshaei, M., Zhu, Q., Alpcan, T., Bacşar, T., Hubaux, J.P.: Game theory meets network security and privacy. ACM Computing Surveys 45(3), 25:1–25:39 (July 2013)
15. Nakamoto, S.: Bitcoin: A peer-to-peer electronic cash system. `http://bitcoin.org/bitcoin.pdf` (2008)
16. Plohmann, D., Gerhards-Padilla, E.: Case study of the miner botnet. In: Proceedings of the 4th International Conference on Cyber Conflict (CYCON). pp. 345–360 (2012)
17. Schechter, S., Smith, M.: How much security is enough to stop a thief? In: Wright, R. (ed.) Financial Cryptography, Lecture Notes in Computer Science, vol. 2742, pp. 122–137. Springer Berlin Heidelberg (2003)
18. Schelling, T.: The Strategy of Conflict. Oxford University Press, Oxford, UK (1965)
19. Spyridopoulos, T., Karanikas, G., Tryfonas, T., Oikonomou, G.: A game theoretic defence framework against DoS/DDoS cyber attacks. Computers & Security 38, 39–50 (Oct 2013)
20. Vasek, M., Thornton, M., Moore, T.: Empirical analysis of denial-of-service attacks in the bitcoin ecosystem. In: Proceedings of the 1st Workshop on Bitcoin Research (BITCOIN). pp. 57–71. Lecture Notes in Computer Science, Springer (March 2014)
21. Wu, Q., Shiva, S., Roy, S., Ellis, C., Datla, V.: On modeling and simulation of game theory-based defense mechanisms against DoS and DDoS attacks. In: Proceedings of the 2010 Spring Simulation Multiconference. pp. 159:1–159:8 (2010)