Game-Theoretic Analysis of DDoS Attacks Against Bitcoin Mining Pools

Benjamin Johnson[1], Aron Laszka[2], Jens Grossklags[3(✉)], Marie Vasek[4], and Tyler Moore[4]

[1] University of California, Berkeley, CA, USA
johnsonb@ischool.berkeley.edu
[2] Budapest University of Technology and Economics, Budapest, Hungary
laszka@crysys.hu
[3] The Pennsylvania State University, State College, PA, USA
jensg@ist.psu.edu
[4] Southern Methodist University, Dallas, TX, USA
{mvasek,tylerm}@smu.edu

Abstract. One of the unique features of the digital currency Bitcoin is that new cash is introduced by so-called *miners* carrying out resource-intensive proof-of-work operations. To increase their chances of obtaining freshly minted bitcoins, miners typically join *pools* to collaborate on the computations. However, intense competition among mining pools has recently manifested in two ways. Miners may invest in additional computing resources to increase the likelihood of winning the next mining race. But, at times, a more sinister tactic is also employed: a mining pool may trigger a costly distributed denial-of-service (DDoS) attack to lower the expected success outlook of a competing mining pool. We explore the trade-off between these strategies with a series of game-theoretical models of competition between two pools of varying sizes. We consider differences in costs of investment and attack, as well as uncertainty over whether a DDoS attack will succeed. By characterizing the game's equilibria, we can draw a number of conclusions. In particular, we find that pools have a greater incentive to attack large pools than small ones. We also observe that larger mining pools have a greater incentive to attack than smaller ones.

Keywords: Game theory · Bitcoin · Internet · Security · DDoS

1 Introduction

Bitcoin is a decentralized digital currency that first became operational in 2009 [1]. While cryptographically protected digital currencies have been around for decades [2], none has received the attention or experienced the same rise in adoption as Bitcoin [3].

There are many factors that contribute to the success of a currency. Most currencies are tightly associated with a particular country, and are influenced

R. Böhme et al. (Eds.): FC 2014 Workshops, LNCS 8438, pp. 72–86, 2014.
DOI: 10.1007/978-3-662-44774-1_6

by decisions regarding economic factors and political leadership. At the same time, internal stakeholders and external trade partners benefit from the adoption and maintenance of a stable currency. Wider adoption enables positive network effects, e.g., by enabling exchange of goods beyond the scope of a traditional barter community. However, currencies remain in competition with each other, and new currencies might gain a foothold if they offer comparative advantages to a certain set of stakeholders [4].[1]

One reason why Bitcoin has attracted enthusiastic backers is that its design creates opportunities for participants to shape its future and to profit from its success. The artificially constrained money supply helps drive up the exchange rate over time, rewarding those who have invested in bitcoins. Most importantly, new bitcoins are given as rewards to the miner who finds the solution to a complex mathematical problem. However, this also means that new entrants in the market for Bitcoin mining impose negative externalities on other contributors. Each new miner who contributes to Bitcoin automatically lowers the value of the relative contributions of all other miners.

Miners respond in two primary ways to increase their output during the quest to earn another bundle of bitcoins. First, they form associations with other contributors in *mining pools*. Second, they may invest in additional computing resources. For example, the increasing value of Bitcoin has also created a market for specialized hardware. At the same time, botnets have been used to increase the output of mining pools that control these illegally acquired resources. In the end, the most powerful mining pool is the most likely to win the next race.

There is one caveat to this relatively straightforward process. More recently, attacks hampering the effectiveness of mining pools have been observed. Distributed Denial of Service Attacks (DDoS) frequently target mining pools in order to disrupt their operations (e.g., the distribution and submission of delegated tasks). There are two primary objectives that attackers are following when facilitating DDoS attacks on mining pools. First, the operations at competing mining pools are slowed down which might give a decisive (but unfair) advantage in the race for the next bundle of bitcoins. Second, individual miners might become discouraged and decide to leave "unreliable" mining pools as the result of these attacks.[2]

[1] Rules for currency competition may differ by country. For example, in the United States the following rules are of importance. United States money, as identified by the U.S. Code, when tendered to a creditor always legally satisfies a *debt* to the extent of the amount tendered. However, no federal law mandates that a person or an organization must accept United States money as payment for *goods or services not yet provided*. That is, a business might specify a particular currency and therefore increase competition between currencies.

[2] Other attack motivations might include the facilitation of other cybercriminal activities, e.g., using DDoS as a means to extract payments from a mining pool as part of an extortion ploy [5]. Attacks might also be indicative of non-financial objectives, e.g., the earning of reputation in the attacker community or general disagreement with the goals and objectives of the Bitcoin community.

Mining pools have been sporadically targeted by DDoS attacks since 2011. According to an empirical analysis of Bitcoin-related DDoS attacks [6], mining pools are the second-most frequently targeted Bitcoin service after currency exchanges. Of 49 mining pools, 12 experienced DDoS attacks, often repeatedly. At least one mining pool, Altcoin.pw, appears to have shut down due to repeated DDoS attacks.

Our study addresses the trade-off between two different investment dimensions in the context of Bitcoin creation: construction and destruction. Under the construction paradigm, a mining pool may invest in additional computing resources to increase the likelihood of winning the next race. Under the destruction focus, a mining pool may trigger a costly DDoS attack to lower the expected success outlook of a competing mining pool.

We approach the study of this trade-off by developing a series of game-theoretical models. We begin our analysis with a simple model that presents a binary choice between investment and DDoS attack. Subsequently, we expand this baseline model to account for costs and the possibility of attack failure. Our goal is to give the reader initially an intuitive understanding about the impact of the different investment choices. With increasing model complexity, we aim for a heightened degree of realism regarding actual investment decisions.

Our work is important because it contributes to a greater understanding of the inherent risks of the Bitcoin economy. Due to its decentralized nature, international focus and lack of regulation, the existing competing and misaligned interests prevalent in the Bitcoin community can frequently lead to undesirable outcomes. For example, many Bitcoin currency exchanges have a short survival time, often leaving their customers in the lurch [7]. The scenario we study becomes an increasingly central concern to Bitcoin mining pools. With accelerating upfront investment costs to compete in the Bitcoin mining race, the associated risks are ballooning as well, e.g., interference with the mining operations becomes more costly. Responding to such threats requires a good understanding of the economic impact of attacks and potential countermeasures.

Our presentation proceeds as follows. In Sect. 2, we briefly discuss related work with a focus on theoretical research. In Sect. 3, we develop and analyze a series of game-theoretical models. We discuss the practical implications of these analyses and conclude in Sect. 4.

2 Related Work

2.1 Economics of Security Decision-Making

Our model is concerned with DDoS attacks as a strategic choice impacting the Bitcoin mining race. As such, we focus in our review on research in which adversarial interests are the subject of economic models. However, relatively little work has addressed the strategic choices of attackers and cybercriminals. Fultz and Grossklags model strategic attackers and the competition between those attackers [8]. In their model, attackers and defenders have to be cognizant of inherent interdependencies that shape the impact of offensive and defensive actions [9–11].

Similarly, Clark and Konrad present a game-theoretic model with one defender and one attacker. The defending player has to successfully protect multiple nodes while the attacker must merely compromise a single point [12]. Cavusoglu *et al.* [13] analyze the decision-making problem of a firm when attack probabilities are externally given compared to a scenario when the attacker is explicitly modeled as a strategic player in a game-theoretic framework.

Cremonini and Nizovtsev compare attacker decisions under different scenarios of information availability regarding defensive strength [14]. Schechter and Smith [15] draw upon the economics of crime literature to construct a model of attackers in the computer security context [16]. They derive the penalties and probabilities of enforcement that will deter an attacker who acts as an utility optimizer evaluating the risks and rewards of committing an offense.

Several surveys have summarized the achievements in this area [17–19].

2.2 Economics of DDoS

Research on the economics of DDoS attacks has focused on the organization of an effective defense [20–22]. For example, Liu et al. develop a game-theoretic model of DDoS attacker-defender interactions, and conduct a network simulation study which utilizes their model to infer DDoS attack strategies [20].

More closely related to our work is a paper by Li et al. [23]. They model the incentives of a botnet herder to maintain a zombie network for the primary purpose of renting a sufficiently large subset to a DDoS attacker. They 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). Complementary to this work, Christin et al. investigate the incentives of a group of defenders when they face the threat of being absorbed into a botnet, e.g., for the purpose of a DDoS attack [24]. Their model shows how the bounded rationality of defenders can contribute to lower defensive investments and a higher risk of security compromise.

We are unaware of any economic research that investigates the potential impact of DDoS attacks on the Bitcoin economy.

2.3 Incentive Modeling of the Bitcoin Economy

In this subsection, we briefly report on research studies that investigate the stability of Bitcoin to economically-driven attacks. We do not review research on the robustness of the cryptographic underpinnings of Bitcoin.

Kroll et al. study the stability of Bitcoin mining if an outsider has motivation to destroy the currency [25]. More specifically, their "Goldfinger" attack compares on a high level the collective benefit of Bitcoin mining with some externally given incentive to destroy the economy altogether. They also study the likelihood of deviations from the consensus process of Bitcoin mining.

Similarly, Barber et al. perform an in-depth investigation of the success of Bitcoin, and study the characteristics of a "doomsday" attack in which the

complete transaction history would be invalidated by an adversary with vastly superior computing power [3]. They also investigate a number of other potential weaknesses, and propose improvements to the Bitcoin protocol.

Babaioff et al. show that, as the Bitcoin protocol is currently defined, it does not provide incentives for nodes to broadcast transactions; in fact, it provides strong disincentives [26]. However, the Bitcoin economy seems to be – at least in this respect – working well in practice. The authors propose a solution for this potential problem, which is based on augmenting the Bitcoin protocol with a scheme for rewarding information propagation.

3 Game-Theoretic Model and Analysis

Our modeling approach focuses on the incentives of Bitcoin mining pool operators to initiate distributed denial of service attacks against other mining pools. Toward this end, we begin our analysis with a very simple model that presents a binary choice between investment and attack. Subsequently, we expand the baseline model to account for the possibility of attack failure, and then to consider linear investment and attack costs.

In each model, we focus on exactly two players – a big player B and a small player S. By the size comparison, we simply mean that B has more computational power to mine bitcoins than S. A third entity R represents the rest of the Bitcoin mining market. R behaves heuristically and thus is not a player in a game-theoretical sense. In equations, we overload the notation B, S, and R to represent the value of the respective player's computing power.

Each player's decision space involves a binary choice of investment – either to invest in additional computing power, or to initiate a DDoS attack against the other strategic player. The outcome of each player's decision is realized over a time scale that is long enough so that payoffs to pools in bitcoins are realized according to the mining probabilities, but short enough so that reaching an approximate equilibrium in the relative computational power of mining pools is a reasonable assumption.

3.1 Baseline Model

We assume that the Bitcoin mining market increases computational power over the game's time scale at a fixed rate ε; and that the market is at an equilibrium with respect to each player's relative computing power. Each player's base strategy is to maintain the market equilibrium by investing in computation to keep up with the market. Each player's alternative strategy is to use those resources that would have been used for increased computation to initiate a DDoS attack against the other strategic player.

In the baseline model, we assume that DDoS attacks are 100 % effective, so that a player who is subject to the attack cannot mine any Bitcoins for the duration of the game's time scale. Secondly, in the baseline model, we assume

that the costs to invest or initiate an attack are negligible relative to the overall Bitcoin revenue, so that they do not factor into the players' strategic decisions.

The payoff for each player is determined by the expected value of the fraction of Bitcoins that they mine. If both players use the base strategy to keep up with the market, then the payoff of player S is

$$\frac{S(1+\varepsilon)}{(B+S+R)(1+\varepsilon)} = \frac{S}{B+S+R};$$

similarly, the payoff for player B is

$$\frac{B}{B+S+R}.$$

If both players initiate DDoS attacks against each other, then they each receive nothing. If player S initiates a DDoS attack against player B, while B keeps up with the market, then B receives nothing, and S receives

$$\frac{S}{S+R(1+\varepsilon)}.$$

These consequences are symmetric with respect to S and B.

The full payoff matrix for each player is summarized in Table 1. From this, we derive each players' best responses to each of the other player's strategies. Then we use best response conditions to classify the game's Nash equilibria. Finally, we provide numerical illustrations for the game's equilibria and analyze the corresponding implications.

Table 1. Payoff matrix for B, S

		Player B	
		Computation	DDoS
Player S	Computation	$\frac{B}{B+S+R}, \frac{S}{B+S+R}$	$\frac{B}{B+R(1+\varepsilon)}, 0$
	DDoS	$0, \frac{S}{S+R(1+\varepsilon)}$	$0, 0$

Best-Response Strategies. If player S invests in DDoS, then investing in DDoS and investing in computing are both best responses for player B, since they both yield a payoff of 0. On the other hand, if player S invests in computing, then investing in DDoS is a unique best response for player B if

$$\frac{B}{B+R(1+\varepsilon)} > \frac{B}{(B+S+R)};$$

which reduces to

$$R\varepsilon < S. \tag{1}$$

Both DDoS and computing are best responses if

$$R\varepsilon = S; \tag{2}$$

and computing is a unique best response otherwise. The best responses of player S analogous, with the constants B and S swapped.

Equilibria

- First, both players investing in DDoS is always a Nash equilibrium. However, this is only a weak equilibrium, as both players are indifferent to their strategy choices.
- Second, both players investing in computing is an equilibrium if

$$S \leq R\varepsilon \tag{3}$$

and

$$B \leq R\varepsilon. \tag{4}$$

Furthermore, the equilibrium is strict if both inequalities are strict.
- Finally, if only one of the above inequalities holds, then there is an equilibrium in which the player whose inequality does not hold invests in DDoS, while the other player invests in computing. This is again a weak equilibrium, since the latter player is indifferent to her strategy.

Numerical Illustrations. Figure 1 shows features of the Nash equilibria for various values of B and S. Figure 1a divides the parameter space based on the set of equilibrium profiles. Figure 1b shows the payoff of player B as a function of the relative sizes of B and S, where the average payoff is taken for regions having multiple equilibria. The average payoffs of players B and S (for a fixed S) are shown as a function of B by Fig. 1c.

From Fig. 1a, we see immediately that it is always a weak equilibrium for each player to DDoS the other. This happens because, with perfect effectiveness of DDoS, the player being attacked loses all incentives related to their strategic choice, and thus can choose an arbitrary strategy. We extend the model in the next section to incorporate imperfect DDoS, which alleviates this phenomenon. From the same figure, we also see that if either player becomes much larger than the market growth rate, there is no incentive to mutually cooperate. In these regions, one of the players always has a greater incentive to DDoS if her opponent invests in computation. The slant of the dividing lines also shows that the tendency to avoid cooperation is slightly affected by a player's own size. Figure 1b shows that in this model, the large player fares extremely poorly against a small player if her size becomes too large relative to the market growth rate.

3.2 Baseline Model with Imperfect DDoS

In the first extension of our baseline model, we assume that DDoS attacks are successful only with fixed probability $1-\sigma$. For numerical illustrations, we take σ

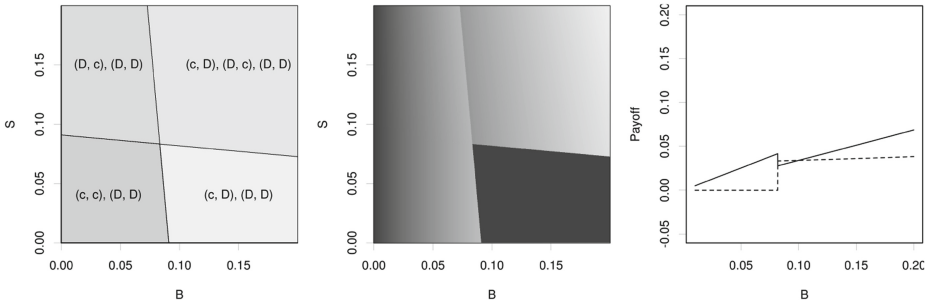

(a) Equilibrium strategy profiles for players (B, S) as a function of the players' sizes. The letters c and D abbreviate computation and DDoS, respectively.

(b) Equilibrium payoff of player B (lighter shades represent higher payoffs). Where there are multiple equilibria, the figure shows the average payoff.

(c) Average equilibrium payoffs of players B (solid) and S (dotted) as a function of B, with $S = 0.1$.

Fig. 1. Equilibria for various values of B and S. The increase in computational power is $\varepsilon = 0.1$.

Table 2. Payoff matrix for B, S with imperfect DDoS

		B	
		Computation	DDoS
S	Computation	$\frac{B}{B+S+R}, \frac{S}{B+S+R}$	$\frac{B}{B+(\sigma S+R)(1+\varepsilon)}, \frac{\sigma S(1+\varepsilon)}{B+(\sigma S+R)(1+\varepsilon)}$
	DDoS	$\frac{\sigma B(1+\varepsilon)}{(\sigma B+R)(1+\varepsilon)+S}, \frac{S}{(\sigma B+R)(1+\varepsilon)+S}$	$\frac{\sigma B}{\sigma(B+S)+R(1+\varepsilon)}, \frac{\sigma S}{\sigma(B+S)+R(1+\varepsilon)}$

to be 0.2. The new payoffs (with arbitrary σ) for players B and S are summarized in Table 2.

Best-Response Strategies. If player S invests in computation, then investing in computation is a best response for player B if

$$\frac{B}{B+S+R} \geq \frac{B}{B+(\sigma S+R)(1+\varepsilon)},$$

which reduces to

$$S \leq \frac{\varepsilon R}{1-\sigma(1+\varepsilon)}; \tag{5}$$

and investing in DDoS is a best response if

$$S \geq \frac{\varepsilon R}{1-\sigma(1+\varepsilon)}. \tag{6}$$

If player S initiates a DDoS attack, then investing in computation is a best response for player B if

$$\frac{\sigma B(1+\varepsilon)}{(\sigma B+R)(1+\varepsilon)+S} \geq \frac{\sigma B}{\sigma(B+S)+R(1+\varepsilon)},$$

which reduces to

$$S \leq \frac{\varepsilon R}{1-\sigma-\frac{\varepsilon}{1+\varepsilon}}; \tag{7}$$

and investing in DDoS is a best response if

$$S \geq \frac{\varepsilon R}{1-\sigma-\frac{\varepsilon}{1+\varepsilon}}. \tag{8}$$

Equilibria. The game's equilibria depend on the sizes of B and S compared to the quantities $\frac{\varepsilon R}{1-\sigma(1+\varepsilon)}$ and $\frac{\varepsilon R}{1-\sigma-\frac{\varepsilon}{1+\varepsilon}}$. Note that we would expect the first quantity to be smaller, because we typically have $\sigma < \frac{1}{1+\varepsilon}$. Concretely, for example, this desired relation holds when the growth rate ε is less than 100 % and the DDoS failure rate σ is at most 50 %.

- First, both players investing in DDoS is a Nash equilibrium whenever

$$B, S \geq \frac{\varepsilon R}{1-\sigma-\frac{\varepsilon}{1+\varepsilon}} \tag{9}$$

  and the equilibrium is strict whenever the inequality is strict.
- Second, both players investing in computing is an equilibrium if

$$B, S \leq \frac{\varepsilon R}{1-\sigma(1+\varepsilon)} \tag{10}$$

  and again the equilibrium is strict if the inequality is strict.
- Third, there exists an equilibrium in which S initiates a DDoS attack and B invests in computation whenever

$$B \geq \frac{\varepsilon R}{1-\sigma(1+\varepsilon)} \tag{11}$$

  and

$$S \leq \frac{\varepsilon R}{1-\sigma-\frac{\varepsilon}{1+\varepsilon}}. \tag{12}$$

- Finally, there is a sub-case of the previous condition in which B can initiate a DDoS attack while S invests in computation, if

$$\frac{\varepsilon R}{1-\sigma(1+\varepsilon)} \leq B, S \leq \frac{\varepsilon R}{1-\sigma-\frac{\varepsilon}{1+\varepsilon}}. \tag{13}$$

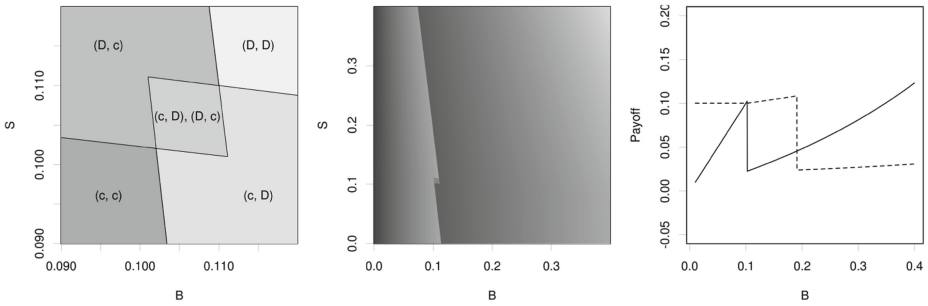

(a) Equilibrium strategy profiles for players (B, S) as a function of the players' sizes. The letters c and D abbreviate computation and DDoS, respectively.

(b) Equilibrium payoff of player B (lighter shades represent higher payoffs). Where there are multiple equilibria, the figure shows the average payoff.

(c) Equilibrium payoff of players B (solid) and S (dotted) as a function of B for $S = 0.1$.

Fig. 2. Equilibria for various values of B and S. The increase in computational power is $\varepsilon = 0.1$, and the success probability of DDoS is $1 - \sigma = 0.8$.

Numerical Illustration. Figure 2, illustrates features of the equilibria for the baseline model with imperfect DDoS. Figure 2a divides the parameter space based on the set of equilibrium profiles. Figure 2b shows the payoff of player B as a function of the relative sizes of B and S; and Fig. 2c shows the payoff of players B and S (for a fixed S) as a function of B.

From Fig. 2a, we see that, (compared to the baseline model) there is no longer a weak equilibrium in which each player initiates a DDoS attack against the other; and in most parameter configurations, there is now a unique equilibrium. For each player, this unique equilibrium strategy is primarily determined by her opponent's computational power. Once the opponent reaches a given threshold, it is in the player's best interest to DDoS that opponent. The slanted nature of the equilibrium-dividing lines shows that a player's equilibrium strategy is also determined to a weaker degree by her own computational power, with larger players having slightly more incentive to attack. Finally, there is a region for players of medium and comparable sizes, in which the game has two competing equilibria. The strategic dynamic in this region is similar to the classical game of *battle of the sexes*.

3.3 Baseline Model with Imperfect DDoS and Linear Costs

The third extension of our baseline model combines the features of imperfect DDoS attacks and linear costs for player investment choices. Here we assume that the cost of an investment to keep up with the mining market is proportional to the size of the investing player, and that the cost to initiate a DDoS attack is proportional to the size of the player who is being attacked.

If S invests in computation, she incurs a cost of γS; while if S initiates a DDoS attack against player B, it results in a cost of λB. Other things being equal, we suppose that a DDoS attack should cost less than an investment in computation, so for our numerical illustrations, we choose an assignment with $\lambda < \gamma$. The resulting payoffs for players B and S (for arbitrary γ and λ) are summarized in Tables 3 and 4.

Table 3. Payoff matrix for B with imperfect DDoS and linear costs

		B	
		Computation	DDoS
S	Computation	$\frac{B}{B+S+R} - \gamma B$	$\frac{B}{B+(\sigma S+R)(1+\varepsilon)} - \lambda S,$
	DDoS	$\frac{\sigma B(1+\varepsilon)}{(\sigma B+R)(1+\varepsilon)+S} - \gamma B$	$\frac{\sigma B}{\sigma(B+S)+R(1+\varepsilon)} - \lambda S$

Table 4. Payoff matrix for S with imperfect DDoS and linear costs

		B	
		Computation	DDoS
S	Computation	$\frac{S}{B+S+R} - \gamma S$	$\frac{\sigma S(1+\varepsilon)}{B+(\sigma S+R)(1+\varepsilon)} - \gamma S$
	DDoS	$\frac{S}{(\sigma B+R)(1+\varepsilon)+S} - \lambda B$	$\frac{\sigma S}{\sigma(B+S)+R(1+\varepsilon)} - \lambda B$

Best-Response Strategies. If player S invests in computation, then investing in computation is a best response for player B if

$$\frac{B}{B+S+R} - \gamma B \geq \frac{B}{B+(\sigma S+R)(1+\varepsilon)} - \lambda S; \tag{14}$$

and investing in DDoS is a best response if

$$\frac{B}{B+S+R} - \gamma B \leq \frac{B}{B+(\sigma S+R)(1+\varepsilon)} - \lambda S. \tag{15}$$

If player S initiates a DDoS attack, then investing in computation is a best response for player B if

$$\frac{\sigma B(1+\varepsilon)}{(\sigma B+R)(1+\varepsilon)+S} - \gamma B \geq \frac{\sigma B}{\sigma(B+S)+R(1+\varepsilon)} - \lambda S; \tag{16}$$

and investing in DDoS is a best response if

$$\frac{\sigma B(1+\varepsilon)}{(\sigma B+R)(1+\varepsilon)+S} - \gamma B \leq \frac{\sigma B}{\sigma(B+S)+R(1+\varepsilon)} - \lambda S. \tag{17}$$

Equilibria

– First, both players initiating DDoS attacks is a Nash equilibrium whenever

$$\frac{B}{B+S+R} - \gamma B \geq \frac{B}{B+(\sigma S+R)(1+\varepsilon)} - \lambda S \tag{18}$$

and

$$\frac{S}{B+S+R} - \gamma S \geq \frac{S}{(\sigma B+R)(1+\varepsilon)+S} - \lambda B. \tag{19}$$

– Second, both players investing in computing is an equilibrium if

$$\frac{\sigma B(1+\varepsilon)}{(\sigma B+R)(1+\varepsilon)+S} - \gamma B \leq \frac{\sigma B}{\sigma(B+S)+R(1+\varepsilon)} - \lambda S \tag{20}$$

and

$$\frac{\sigma S(1+\varepsilon)}{B+(\sigma S+R)(1+\varepsilon)} - \gamma S \leq \frac{\sigma S}{\sigma(B+S)+R(1+\varepsilon)} - \lambda B. \tag{21}$$

– Third, an equilibrium in which S conducts a DDoS attack against B while B invests in computation may occur when

$$\frac{\sigma B(1+\varepsilon)}{(\sigma B+R)(1+\varepsilon)+S} - \gamma B \leq \frac{\sigma B}{\sigma(B+S)+R(1+\varepsilon)} - \lambda S \tag{22}$$

and

$$\frac{S}{B+S+R} - \gamma S \leq \frac{S}{(\sigma B+R)(1+\varepsilon)+S} - \lambda B. \tag{23}$$

– Finally, there can be an equilibrium in which B conducts a DDoS attack against S while S invests in computation whenever the roles of B and S are interchanged in the two inequalities from the previous case.

Numerical Illustration. Figure 3 shows features of the Nash equilibria for various values of B and S. Figure 3a divides the parameter space based on the set of equilibrium profiles. Figure 3b shows the payoff of player B as a function of the relative sizes of B and S; and Fig. 3c shows the payoff of players B and S (for a fixed S) as a function of B.

The addition of costs to the model keeps the smallest players from participating in DDoS attacks, as they are best served by investing in their own computational prowess. Aside from this, the dynamics of the equilibrium strategies are largely similar to the model without costs. Namely, players are still incentivized to attack large players, and slightly more so if they are larger themselves. There still remains a small region for midsize players in which either player can attack the other; and with the possible exception of an extremely large player, the payoffs are generally higher for a player whose size lies just below the threshold for being attacked.

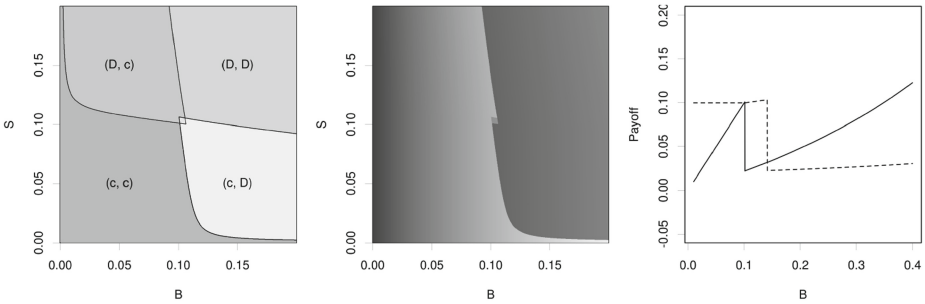

(a) Equilibrium strategy profiles for players (B, S) as a function of the players' sizes. The letters c and D abbreviate computation and DDoS, respectively.

(b) Equilibrium payoff of player B (lighter shades represent higher payoffs). Where there are multiple equilibria, the figure shows the average payoff.

(c) Equilibrium payoff of players B (solid) and S (dotted) as a function of B for $S = 0.1$.

Fig. 3. Equilibria for various values of B and S. The increase in computational power is $\varepsilon = 0.1$, the success probability of DDoS is $1 - \sigma = 0.8$, and the linear cost factors for investing into computation and DDoS are $\gamma = 0.002$ and $\lambda = 0.001$.

4 Conclusions and Future Work

We set out in this work to understand the motivation behind recent DDoS attacks against Bitcoin mining pools. To do this, we analyzed a series of game-theoretical models involving two mining pools with different sizes. Several fundamental dynamics of this game were common to all models and seem well-motivated in the context of Bitcoin. First, we saw that there is a greater incentive to attack a larger mining pool than a smaller one. This finding is intuitive because each pool battles for the reward; and eliminating the largest mining pool has the greatest impact on the chances of the remaining mining pools to win. It is also consistent with what has been observed empirically: 63 % of large mining pools have experienced DDoS attacks, compared to just 17 % of small ones [6]. Second, we observed that the larger mining pool has a slightly greater incentive to attack than the smaller mining pool. This dynamic arises because a larger mining pool has a smaller relative competitor base, and eliminating a competitor from a small base yields more benefit than eliminating one from a larger base. Finally, there is a size threshold such that mining pools larger than this threshold are subject to economically-motivated attacks; and pools smaller than the threshold are not. Furthermore, players whose sizes are just below this threshold tend to receive the highest payoffs.

From our modeling extensions we found additional insights. First, if attacks can be mitigated, then the size threshold for a mining pool to be safe from DDoS increases. That is, the market will tolerate (without attempting an attack) progressively larger pools as attacks become less effective. Second, the prevalence

of costs can keep smaller players out of the DDoS market, but these do not change the core dynamics for mid-size and large mining pools.

There are many extensions to pursue in future work. A more direct economic approach to the cost dimension would have each player optimize their own investment costs relative to their current size. A player's choice of whether to initiate a DDoS attack would depend on the solution to two investment optimization problems. This extension would improve realism and reduce the game's exogenous parameters at the expense of additional model complexity. Another way to extend the model would be to give DDoS attacks a variable cost constraining their effectiveness. Finally, our work considers the incentives of mining pools as a whole, but in reality most pools consist of heterogeneous individuals who have a choice to change pools. By expanding our game to an iterated version in which individual players could switch mining pools between rounds, we might gain further insights into the strategies we see in today's Bitcoin mining market.

Acknowledgements. This research was partly supported by the Penn State Institute for CyberScience, CyLab at Carnegie Mellon under grant DAAD19-02-1-0389 from the Army Research Office, and the National Science Foundation under ITR award CCF-0424422 (TRUST). We also thank the reviewers for their comments on an earlier draft of the paper.

References

1. Nakamoto, S.: Bitcoin: a peer-to-peer electronic cash system. http://bitcoin.org/bitcoin.pdf (2008)
2. Chaum, D., Fiat, A., Naor, M.: Untraceable electronic cash. In: Goldwasser, S. (ed.) CRYPTO 1988. LNCS, vol. 403, pp. 319–327. Springer, Heidelberg (1990)
3. Barber, S., Boyen, X., Shi, E., Uzun, E.: Bitter to better — how to make Bitcoin a better currency. In: Keromytis, A.D. (ed.) FC 2012. LNCS, vol. 7397, pp. 399–414. Springer, Heidelberg (2012)
4. Dowd, K., Greenaway, D.: Currency competition, network externalities and switching costs: towards an alternative view of optimum currency areas. Econ. J. **103**(420), 1180–1189 (1993)
5. 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)
6. Vasek, M., Thornton, M., Moore, T.: Empirical analysis of Denial-of-Service attacks in the Bitcoin ecosystem. In: Böhme, R., Brenner, M., Moore, T., Smith, M. (eds.) FC 2014 Workshops. LNCS, vol. 8438, pp. 57–71. Springer, Heidelberg (2014)
7. Moore, T., Christin, N.: Beware the middleman: empirical analysis of Bitcoin-exchange risk. In: Sadeghi, A.-R. (ed.) FC 2013. LNCS, vol. 7859, pp. 25–33. Springer, Heidelberg (2013)
8. Fultz, N., Grossklags, J.: Blue versus red: towards a model of distributed security attacks. In: Dingledine, R., Golle, P. (eds.) FC 2009. LNCS, vol. 5628, pp. 167–183. Springer, Heidelberg (2009)
9. 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), Beijing, China, April 2008, pp. 209–218 (2008)

10. Grossklags, J., Johnson, B., Christin, N.: When information improves information security. In: Sion, R. (ed.) FC 2010. LNCS, vol. 6052, pp. 416–423. Springer, Heidelberg (2010)
11. Varian, H.: System reliability and free riding. In: Camp, L., Lewis, S. (eds.) Economics of Information Security. Advances in Information Security, vol. 12, pp. 1–15. Kluwer, Dordrecht (2004)
12. Clark, D., Konrad, K.: Asymmetric conflict: weakest link against best shot. J. Conflict Resolut. **51**(3), 457–469 (2007)
13. Cavusoglu, H., Raghunathan, S., Yue, W.: Decision-theoretic and game-theoretic approaches to IT security investment. J. Manag. Inf. Syst. **25**(2), 281–304 (2008)
14. Cremonini, M., Nizovtsev, D.: Understanding and influencing attackers' decisions: Implications for security investment strategies. In: Proceedings of the Fifth Annual Workshop on Economics and Information Security (WEIS), Cambridge, UK, June 2006
15. Schechter, S.E., Smith, M.D.: How much security is enough to stop a thief? In: Wright, R.N. (ed.) FC 2003. LNCS, vol. 2742, pp. 122–137. Springer, Heidelberg (2003)
16. Becker, G.: Crime and punishment: an economic approach. J. Polit. Econ. **76**(2), 169–217 (1968)
17. Anderson, R., Moore, T.: The economics of information security. Science **314**(5799), 610–613 (2006)
18. Laszka, A., Felegyhazi, M., Buttyán, L.: A survey of interdependent security games. Technical report CRYSYS-TR-2012-11-15, CrySyS Lab, Budapest University of Technology and Economics (2012)
19. Manshaei, M., Zhu, Q., Alpcan, T., Bacşar, T., Hubaux, J.P.: Game theory meets network security and privacy. ACM Comput. Surv. **45**(3), 25:1–25:39 (2013)
20. Liu, P., Zang, W., Yu, M.: Incentive-based modeling and inference of attacker intent, objectives, and strategies. ACM Trans. Inf. Syst. Secur. **8**(1), 78–118 (2005)
21. Spyridopoulos, T., Karanikas, G., Tryfonas, T., Oikonomou, G.: A game theoretic defence framework against DoS/DDoS cyber attacks. Comput. Secur. **38**, 39–50 (2013)
22. 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)
23. Li, Z., Liao, Q., Striegel, A.: Botnet economics: uncertainty matters. In: Johnson, M. (ed.) Managing Information Risk and the Economics of Security, pp. 245–267. Springer, Heidelberg (2009)
24. 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)
25. Kroll, J., Davey, I., Felten, E.: The economics of Bitcoin mining, or Bitcoin in the presence of adversaries. In: Proceedings of the Twelfth Annual Workshop on Economics and Information Security (WEIS), Washington, DC, June 2013
26. Babaioff, M., Dobzinski, S., Oren, S., Zohar, A.: On Bitcoin and red balloons. In: Proceedings of the 13th ACM Conference on Electronic Commerce (EC), pp. 56–73 (2012)