Cryptoeconomic Systems

SoK: Tools for Game Theoretic Models of Security for Cryptocurrencies

Sarah Azouvi¹, Alexander Hicks²

¹University College London, Protocol Labs, ²University College London

Published on: Nov 26, 2020

License: [Creative Commons Attribution 4.0 International License \(CC-BY 4.0\)](https://creativecommons.org/licenses/by/4.0/)

Abstract

Cryptocurrencies have garnered much attention in recent years, both from the academic community and industry. One interesting aspect of cryptocurrencies is their explicit consideration of incentives at the protocol level, which has motivated a large body of work—yet many open problems still exist, and current systems rarely deal with incentive related problems well. This issue arises due to the gap between cryptography and distributed systems security, which deals with traditional security problems that ignore the explicit consideration of incentives, and game theory, which deals best with situations involving incentives. With this work, we offer a systematization of the work that relates to this problem, considering papers that blend game theory with cryptography or distributed systems. This gives an overview of the available tools, and we look at their (potential) use in practice in the context of existing blockchain-based systems that have been proposed or implemented.

Keywords: cryptocurrency, consensus, game theory, cryptoeconomics, blockchain, security

1. Introduction

Since the deployment of Bitcoin in 2009, cryptocurrencies have garnered much attention from both academia and industry. Many challenges in this area have since been recognized, from privacy and scalability to governance and economics. In particular, the explicit consideration of incentives in the protocol design of cryptocurrencies (or “cryptoeconomics”) has become an important topic.

The importance of economic considerations in security has been acknowledged since work by Anderson [\[1\]\[2\]](#), who recognized that many security failures could be explained by applying established ideas from game theory (GT) and economics. However, the incentives at play tend to be external to the system design, and sometimes implicit, leading to failures when the intended use of systems is misaligned with the incentives of users.

Cryptocurrencies, on the other hand, explicitly define some incentives in the design of their system (*e.g.*, in the form of mining rewards), suggesting that they could be properly aligned and avoid traditional failures. Unfortunately, many attacks related to incentives have nonetheless been found for many cryptocurrencies [\[3\]\[4\]\[5\]](#), due to the use of incomplete models. While many papers aim to consider both standard security and game theoretic guarantees, the vast majority end up considering them separately despite their relation in practice.

Here, we consider the ways in which models in cryptography and distributed systems (DS) can explicitly consider game-theoretic properties and incorporate them into a system, looking at requirements based on existing cryptocurrencies.

Methodology

As we are covering a topic that incorporates many different fields, coming up with an extensive list of papers would have been quite challenging, and would lead to an output of much greater length. In order to pick a representative subset of papers, we started by looking at existing surveys on the topic of game theory and security [\[6\]](#)[\[7\]](#)[\[8\]](#)[\[9\]](#)[\[10\]](#)[\[11\]](#), as well as specific book chapters on the topic *e.g.*, Chapter 8 in the book by Nisan *et al.* [\[12\]](#). We then looked at work published in popular cryptography venues (*e.g.*, IACR conferences), security conferences (*e.g.*, FC), as well as distributed system venues (*e.g.*, PODC) and interdisciplinary venues (*e.g.*, WEIS, ACM Economics and Computation, P2PECON, GameSec), looking specifically for papers that cover both game theory and security.

Different papers sometimes present different definitions for similar concepts, so we do not always include all these definitions. We also omit to present work on security and game theory that does not directly relate to what we discuss; *e.g.*, the body of work by Tambe *et al.* [\[13\]](#) about physical security and allocation of limited security resources, or by Grossklags *et al.* [\[14\]](#) about security investments. Instead, we focus on some specific models of interest that we think are more appropriate, and match these with problems in the security of cryptocurrencies.

Our contributions

Our goal is to give an overview of the intersection of the three fields that are essential to the design of cryptocurrencies: cryptography, distributed systems, and game theory. Our contribution is an analysis of existing work that proposes solutions to this problem. Our analysis highlights new concepts introduced by these papers, as well as deficiencies. We do this in the context of security requirements that we formulate, arguing that they address deficiencies in existing security models that fail to cover all aspects of a decentralized cryptocurrency. We also discuss open challenges and how they could be addressed.

[Section 3](#) introduces game theory and cryptocurrencies, and discusses security in the context of a decentralized system. In [Section 4](#) we then look at the intersection of cryptography and game theory, followed by the intersection of DS and game theory in [Section 5](#). We then look at how these results are used in [Section 6](#), where we look at

proposed systems and their failures, tracing back to deficiencies identified in the two previous sections. In [Section 7](#) we give a comparative table of some of the concepts presented in this paper and how they could be used in blockchain research, and we discuss the open challenges posed by failures that are observed, as well as how they could be addressed. Finally, [Appendix A](#) collects formal definitions for the concepts presented in the paper.

2. Related work

The work that most closely resembles ours are the previous surveys bridging computer science and game theory [\[6\]\[7\]\[8\]\[9\]\[10\]\[11\]](#). They were of great inspiration for this work, but they are quite outdated (dating back to 2002, 2005, 2007, 2008, 2010) given the recent output of research tied to cryptocurrencies and other blockchain-based systems.

On the topic of blockchains, many Systems of Knowledge (SoK) papers and surveys exist that cover consensus protocols and security [\[15\]\[16\]\[17\]\[18\]\[19\]\[20\]\[21\]](#). These are very different from our work as we present general concepts and definitions related to designing decentralized systems with incentives. In particular, many concepts presented in this paper were not introduced in the context of consensus, but rather in the context of secure multiparty computation (MPC) or other problems tied to distributed systems. Most of the work presented in this SoK does not directly concern blockchains, which is the motivation behind this work.

3. Background

In this section we briefly introduce game theoretic tools that are mentioned throughout the paper such as solutions concepts and mechanism design (MD). For a complete introduction to game theory, the reader is invited to look at any of the books (or other resources) on the topic [\[22\]\[23\]](#). For the sake of exposition we have opted not to cover concepts that are relevant; *e.g.*, correlated equilibria, Pareto efficiency, and the single deviation test, as these do not appear in the papers we mention. We also discuss the interface between game theory and cryptography and security in practice, and briefly introduce distributed systems and cryptocurrencies.

3.1 Game Theory and Mechanism Design

A game is defined by a set of players and a set of actions for each player. A strategy for a player is defined as a function from its local state to actions.

Game theory uses solution concepts in order to predict the outcome of a game, the most well known is the *Nash equilibrium* (NE). A strategy is a Nash equilibrium if, given that all the other players follow this strategy, player i is better off playing it, as well.

It is often unrealistic to assume that players have complete information about the game they are in. A game where players do not always know what has taken place earlier in the game is said to have *imperfect information*. In the case where players do not know the *type* of the other players, which determines their utility function, the game is said to have *incomplete information*.

The players then have a probability distribution over the types of other players. Their beliefs are expressed as conditional probabilities based on the information they have, which they can update using Bayes' theorem when they gain new information, leading games of this form to be called *Bayesian games*.

Players now also think about expected payoffs, so the standard definition of a NE is no longer ideal but we can define a *Bayesian Nash equilibrium* (BNE) analogously by replacing utilities with expected utilities, although we still refer to them simply as utilities.

While Game Theory is typically about understanding the behavior of players in a given game, systems are usually designed and implemented with a goal in mind (*e.g.*, preventing double spending in cryptocurrencies). To achieve this, our goals can be expressed as a social choice function (SCF), a function that, given the preference (or types) of all players, outputs an outcome. For example, in a voting system, given all the ranked preferences of voters, a SCF will choose a candidate.

Once we have a target outcome in mind, the idea is to make sure that the incentives are designed in a way such that selfish players reach this outcome. In some ways, this can be thought of as reversing the basic idea of game theory, designing a game that leads to a specific outcome.

For example, in a voting system, we would like to design a system where (given all the preferences of the players) the one chosen by the SCF is elected, and one way to do this is to incentivize players to report their truthful preferences. A mechanism can also be viewed as a protocol, with the corresponding game being thought of as having the protocol as the recommended strategy, and deviations from the protocol as other possible strategies.

It must be pointed out that doing this in practice is not always easy, as experimental game theory reveals. Gneezy & Rustichini [24] looked at the effects of implementing fines at a nursery in order to reduce the rate at which parents collected their child late. Intuitively, this should motivate parents to arrive on time, but parents instead interpreted the fine as a way of paying for extra childcare and started coming even later. Furthermore, once the fine was removed (as it was counterproductive), the parents behavior did not revert back—so the system had been irreversibly damaged.

A very important result in MD is the revelation principle, which states that any social choice function that can be implemented by any mechanism can be implemented by a direct truthful mechanism. A mechanism is direct if players need only reveal their type/utility function to the designer of the game, and truthful if players' best strategy is to reveal their true type.

3.2 Agents in Game Theory and Security

Both game theory and security deal with the interaction of agents, but they differ noticeably in how they model these agents. Security deals with *adversaries*, agents that aim to circumvent or break a security property of the system. The value that an adversary attaches to their success is not usually given, as security should ideally be robust to any adversary, although they may have computational limitations. Game theory, on the other hand, deals with *rational* (sometimes also called *selfish*) agents that assign a value to their goals, and would rather optimize their payoff than achieve an arbitrary goal, but they typically do not have restrictions (*e.g.*, computational) like a security adversary would.

In practice, this translates to different assumptions being used when formally modelling a game or the security of a system, making it difficult to prove statements involving security and game-theoretic properties. As both deal with different types of agents, a proof will involve complexity from both sides and quickly become hard to manage, leading to a disjoint treatment of both aspects in works that attempt to cover both. There are nonetheless inherent connections as security often uses game-based proofs, although an adversary wins if they have a high enough probability of succeeding in their attack rather than if their utility is high enough. But if we assume that the adversary has a high payoff associated with the success of their attack, then we start to recover game-theoretic intuition.

For proofs based on the simulation (ideal/real world) paradigm, some connections are also evident. The idea behind the simulation paradigm first originated in the context of secure computation. Goldreich *et al.* [25] introduced the idea of bypassing the need for

a trusted third party (*i.e.*, a mediator) in games of incomplete information, by replacing it with a protocol that effectively simulates it such that any information known by players at any step of game is the same as they would have known in an execution of the game involving the trusted third party. This was refined by Micali & Rogaway [26] in terms of ideal and secure function evaluation. The ideal function evaluation corresponds to the evaluation of the function with a trusted third party that receives the private inputs of the parties and evaluates the function before returning the result to each party, achieving the best possible result. The secure function evaluation involves the parties trying to replace the trusted third party with a protocol, which is considered secure if the parties cannot distinguish between the ideal and secure function evaluations, meaning that an adversary is not able to gain anything significant. Simulation has become a powerful tool for cryptographers [27] and Canetti's Universal Composability (UC) Model [28] expands on these ideas to provide a framework for secure composability of protocols.

3.3 Decentralization, incentives and security

Because the security of most decentralized systems, like cryptocurrencies, is linked not only to the security of the protocols, but also to having a majority of participants following the rules, decentralization and incentives have to be considered.

The ideal system does not depend in any way on any single party, which requires it to be decentralized. Troncoso *et al.* [29] give an overview of decentralized systems, defining a decentralized system as “a distributed system in which multiple authorities control different components and no single authority is fully trusted by all others.” This highlights the fact that every component of the system should be decentralized, and in particular, a single authority distributing its own system (or component) is not decentralized. This can be hard to achieve in practice, and the level of decentralization of a system should always be looked at critically. A decentralized system where all important parties are independent but under the jurisdiction of a single government may not truly be decentralized. All these independent parties may also depend on a very few hardware manufacturers (or other service providers).

Incentives are key to achieving an honest majority. Azouvi *et al.* [30] give an overview of the role incentives play in security protocols (including cryptocurrencies), highlighting the fact that achieving guarantees of equilibria on paper may not be meaningful in practice when the wrong assumptions and models are used. What does security mean in this context? Clearly, protocols that are cryptographically secure, and that achieve safety (*i.e.*, the guarantee that nothing bad will happen) and liveness (*i.e.*,

the guarantee that something good will happen) are needed, otherwise nothing else would work. But if the security of the system also depends on achieving a high enough degree of decentralization, more is required. In particular, decentralization relates to the participants and their behavior rather than solely to the protocol. Any decentralized protocol can always run in a centralized manner, so it is not enough to design a system that can be used in a decentralized manner. Rather, the requirement is to design a system that is advantageous to use in decentralized manner.

Doing this naturally requires a better understanding of why users would want to be decentralized rather than try to gain more individual control of the system for themselves, so security is no longer just about the protocol itself, but also about how it can be used and how it is used.

Bitcoin represents an important innovation from classical consensus protocols as it is fully open and decentralized. In the Bitcoin consensus protocol (sometimes called Nakamoto consensus), participants can join and leave as they wish, and Sybils are handled through the use of *Proof-of-Work* (PoW).

The data structure that keeps track of the state of the system in Bitcoin is a chain of chronologically ordered blocks (*i.e.*, the blockchain), with each block containing a list of transactions. To win the right to append a block (and win the block reward), participants compete to solve a computational puzzle (*i.e.*, a PoW). They include in their block the solution to that puzzle, the PoW, such that other players can verify its correctness. This block then initiates a new puzzle to be solved. This process of creating new blocks is called *mining* and participants in this protocol are called *miners*.

The security of Bitcoin relies on a majority of the mining power (*i.e.*, hashing power) in the network following the protocol, whether because they are honest or simply rational. Taking control of half of the computational power of Bitcoin for only an hour has a considerable cost (around US \$670 000 [\[31\]](#) as of May 2019), although it is within reach of potential adversaries. This cost depends on the hash rate of the network (*i.e.*, the cost of mining) and the price of Bitcoin in USD, as once mining is no longer profitable for some miners they are likely to stop mining, reducing the hashing power required to control a majority of the network.

This is one of the reasons why Bitcoin's security is so tightly linked to incentives, as when mining is no longer worthwhile the security of the network decreases. Participation in the network is also rewarded by financial gain (through block rewards and transaction fees). The more participants there are, the harder it is to attack the

network since the cost for mounting a 51% attack (where an adversary takes control of more than half of the computational power) increases. These financial motivations are thus also paramount.

Since Bitcoin's deployment, many alternative cryptocurrencies that similarly rely on a blockchain have emerged. The most popular of these is Ethereum [32], which differs from Bitcoin in that it provides a more complex scripting language, meaning that rather than processing simple transactions, nodes in the system execute a script that allows users to perform multitude of functionalities (so-called *smart contracts*).

4. Cryptography and Game Theory

This section considers work at the intersection of game theory and cryptography. Cryptography usually considers a worst-case adversary, but by relaxing this assumption, it is possible to design protocols that bypass impossibility results or achieve better efficiency than existing ones, while maintaining a realistic adversarial model.

4.1 Cryptography meets Game Theory: Rational Cryptography

Initiated by Dodis, Rabin, and Halevi [33], rational cryptography is a subfield of cryptography that incorporates incentives in cryptographic protocols. In this context, new adversaries and their capabilities must be defined, as well as: 1) how to account for incentives; and 2) how protocols can be proven secure for such adversaries.

First, we note that most of this work [34][35][36][37][38][39][40] focuses on multi-party secret sharing or secure function evaluation. Thus, no monetary incentive is usually considered. As pointed out by Dodis & Rabin [40], in a rational cryptographic context, the utilities of the players are usually dependent on cryptographic considerations such as: correctness (*i.e.*, a player prefers to compute the function correctly), exclusivity (*i.e.*, a player prefers that other players do not learn the value of the function correctly), privacy (*i.e.*, a player does not want to leak information to other players), or voyeurism (*i.e.*, a player wants to learn as much as possible about the other parties).

In addition to the above, other interesting parameters can come to into play in the adversary's utility function. For example, Aumann and Lindell [41] formalized the concept of *covert adversaries* that may deviate from the protocol, but only if they are not caught doing so. As they argue, there are many obvious situations where parties cannot afford the effect of being caught cheating.

Covert adversaries are somehow similar to adding a punishment to the utility function. Rational players do not want to be caught cheating as the punishment decreases their utility. In Aumann & Lindell’s setting, the protocol detects the cheating, but in practice, we need to incentivize participants to do so. Some work considers adding adversarial behavior together with rational adversaries [42]; we consider this further in [Section 5](#).

In terms of equilibria, the solution concepts proposed in these works are often extensions of a Nash Equilibria (NE), introduced in [Section 3](#). For example, Halpern & Teague look for a NE that remains after other NE that are weakly dominated (*i.e.*, at best only as good as others) are removed through iterated deletion, where all dominated strategies are removed at each step [34]. Asharov *et al.* [39] adapt the simulation-based definition to capture game-theoretic notions of (for example) fairness, meaning that one party learns the output of a computation if and only if the other does as well. As their notions are weaker than standard cryptographic definitions, they can be achieved in some settings where impossibility results usually hold in traditional cryptography.

Following the work presented above, Garay *et al.* propose Rational Protocol Design (RPD) [43]. In this setting, they define a game between the designer of the protocol D and the attacker A . The game is parametrized by a multi-party functionality $\mathcal{F}$ and consists of two sequential moves. In the first step, D sends to A the description π of the protocol that honest parties are supposed to execute. In the second step, A chooses a polynomial-time interactive Turing machine (ITM) $\mathbf{Adv}$ to attack the protocol. The game is zero-sum in the original paper, but was later adapted to be a non-zero sum game in the context of Bitcoin [44].

In follow-up work [45], notions of fairness are also considered, and provide a means of comparison between protocols (*i.e.*, which protocol is the fairer). Informally, a protocol π will be at least as fair as another protocol π_0 if the utility of the best adversary A attacking π (*i.e.*, the adversary which maximizes $u_A(\pi, A)$) is no larger than the utility of the best adversary attacking π_0 , except for some negligible quantity.

The solution concept introduced within the RPD framework is ϵ -subgame perfect equilibrium where the parties’ utilities are ϵ close to their best response utilities. When it comes to security, the RPD framework defines the notion of attack-payoff security. Informally, attack-payoff security states that an adversary has no incentive to deviate from the protocol.

Another concept, incentive compatibility, was introduced in a follow-up work of RPD [44]. Here, the definition is slightly different than the definition usually given within mechanism design (MD), where participants achieve the best outcome by revealing their true preferences. Informally, incentive compatibility states that agents gain some utility when participating in the protocol (*i.e.*, they choose to play instead of “staying at home”).

Apart from the recent work of Badertscher, *et al.* [44], rational cryptography does not consider monetary payment.

One important drawback of RPD is that it does not consider the presence of irrational adversaries despite the fact that, in security, we do not always know the motivation of an attacker. RPD uses a relaxed functionality to allow for some defined attacks but this may not cover all attacks, leaving the door open to potential attacks. The UC model does not automatically start accounting for all possible incentives—this is a clear flaw, as we know that only arbitrarily considering incentives leads to failures (*e.g.*, failure of considering outside incentives, or in general “soft” incentives like political and other external incentives [30]).

4.2 Game Theory meets Cryptography: computational games

Rather than starting from a cryptographic setting and incorporating game theoretic notions, as presented above, one can also start from a game theoretic setting and from there move towards cryptographic notions by considering the computational aspects of games. This approach is taken in a body of work by Halpern & Pass that considers *Bayesian machine games*, first introduced in a preprint [46] and that has later appeared in different forms [47][48][49], primarily venues focused on economics rather than security.

A Bayesian machine game (BMG) is defined very similarly to a standard Bayesian game (introduced in Section 3), it only differs in that it considers the complexity (in computation, storage cost, time or otherwise) of actions in the game. This is done by having players pick machines (*e.g.*, a TM or ITM) that will execute their actions, and defining a complexity function for that machine, which the utility takes into account.

A Nash equilibrium for a BMG is expressed in the usual way, but it now takes into account the machine profile rather than a strategy profile. There is, however, an important distinction to make between a standard Nash equilibrium and a Nash equilibrium in machine games, which is that the latter may not always exist. The necessary conditions for the existence of a Nash equilibrium in a machine game are

given by Halpern & Pass [46] to be a finite type space, bounded machines, and a computable game. A follow-up paper by Halpern *et al.* [50] discusses the general question of the existence of a Nash equilibrium for resource-bounded players.

So far, the discussion of computational games has not yet touched on security-related issues, but Halpern & Pass prove an equivalence theorem that relates the idea of universal implementation in a BMG to the standard notion of secure computation in cryptography [51][25]. Intuitively, this goes back to the work of Goldreich, Micali, and Widgerson [25] that first expressed (to the best of our knowledge) the idea of secure computation as the replacement for a mediator in a game that preserves an equilibrium.

A universal implementation corresponds to the idea that a BMG implements a mediator if, whenever a set of players want to truthfully provide their input to the mediator, they also want to run their machine using the same input, preserving the equilibrium and action distribution. There are then multiple equivalence theorems of different strength (up to the information theoretic case), that relate flavors of secure computation to flavors of implementation. The relation is important, as it not only implies that secure computation leads to a form of game-theoretic implementation, but also the reverse. This opens up the option that the guarantees of (some flavor of) secure computation could be achieved by considering the game theory of a problem, although it is not clear whether this process would be more efficient.

BMG have natural applications to known security problems. For example, dealing with covert adversaries as described by Aumann & Lindell [41] (introduced above) can be done by introducing a (two player, for example) mediated game where the honest strategy is to report your input to the mediator and output its reply (with utility $\frac{1}{2}$), and the string *punish* can be output by a player to ensure the other receives payoff 0. Then any secure computation with respect to covert adversaries with deterrent (*i.e.*, the probability of getting caught cheating) $\frac{1}{2}$ is an implementation of the mediator as the expected utility of a cheating player will be $\frac{1}{2} \cdot 1 + \frac{1}{2} \cdot 0 = \frac{1}{2}$, which is the same as that of the honest strategy.

5. Distributed Systems and Game Theory

In this section, we present the work that is at the intersection of game theory and distributed systems, looking at concrete problems that have been well studied. For each case, we will illustrate important concepts and techniques used.

5.1 Algorithmic Mechanism Design

Algorithmic mechanism design (AMD) is concerned with designing games such that self-interested players achieve the game designer's goals, in the same way that distributed systems designers aim, for example, to achieve agreement in the presence of Byzantine players. AMD was first introduced by Nisan & Ronen [52], who proposed that an algorithm designer should ensure that the interests of participants in a distributed setting are best served by behaving correctly (*i.e.*, the algorithm designer should aim for incentive compatibility). The framework of Nisan & Ronen is defined for a centralized computation, but it has been extended to distributed algorithmic mechanism design (DAMD) following work by Feigenbaum *et al.* on cost sharing algorithms for multicast transactions [53]. This led to further developments and applications of DAMD to interdomain routing, web caching, peer-to-peer file sharing, application layer overlay networks, and distributed task allocation, which are summarized in a review by Feigenbaum & Shenker [54].

5.2 Public goods, free riding and hidden actions

Public goods, which are produced at a cost but available to use for free, naturally occur in distributed systems. In a public goods game, players choose to contribute a certain amount, with all contributions being combined and distributed among all players. Naturally, this can lead to players rationally deciding to contribute less to maximize their utility. They may even contribute nothing, which is generally referred to as *free riding*.

Varian first considered modelling the reliability of a system as a public good [55]. The reliability can either depend on the total effort (*i.e.*, the sum of the efforts exerted by the individuals), on the weakest link (*i.e.*, the minimum effort) or on the best shot (*i.e.*, the maximum effort). For example, if there is a wall defending a city, its reliability can depend on the sum of all the work provided by the builders (total effort), on the lowest height (weakest link) or if we consider several walls, on the highest one (best shot). In the case of total effort, the NE corresponds to all players free riding on the player with highest benefit-cost ratio. Moreover, the effort exerted in the NE is always lower than the social optimum (*i.e.*, the best outcome across all players).

Peer-to-peer file sharing also provides an interesting case of a networked system that has faced free riding [56]. As explained by Babaioff *et al.* [57], solutions to this problem could be based on a reputation system, barter, or currency. Another approach, which would not need to keep any long term state information, is to replace *indirect*

reciprocity with *direct reciprocity*. For example, a file in BitTorrent is partitioned into smaller chunks, requiring repeat interactions among peers and enforcing more collaboration between them [58]. In practice, however, this has been shown to not be very effective as it is not robust to strategic agents [59] and induces free riding [60]. There is also the problem of dealing with newcomers, as an adversary can create new identities in order to abuse the system. Analyzing the incentives at play, Feldman *et al.* [61] suggest that penalizing all newcomers may be an effective way of dealing with the problem, as it is not possible to penalize only users abusing the system.

In addition to free-riding, there are many other parameters that a selfish player could abuse in a P2P file sharing system (*e.g.*, when to join or leave, who to connect to, untruthful sharing information, and so on). This is the problem of *hidden actions* (*i.e.*, how peers selfishly behave when their actions are hidden from the rest of the network). In order to analyze the degradation due to hidden actions, Babaioff *et al.* [57] apply the principal-agent framework, due to the similarity of the hidden action problem with that of *moral hazard*. This framework is used in economics when one entity, the *principal*, employs a set of n agents to take actions on its behalf. In order to capture the efficiency of a system in that framework, they define the *Price of Unaccountability* of a technology as the worst ratio between the principal's utility in both the observable-actions case and the hidden actions case. Dealing with observable and hidden actions relates to the transparency of the system, which can be approached from a cryptographic point of view to ensure that agents all see the same set of actions [62]. Another solution, Karma [63], proposes a system for peer-to-peer resource sharing that avoids free riding, based on a combination of reputation system and consensus protocols.

Another important problem in distributed systems where rationality can cause problems is routing [64]. The problem is to find a path that minimizes the latency between a source and a target. One of the difficulties in doing so is that in a decentralized communication network it is not always possible to impose some routing strategy to nodes in order to, for example, regulate the load on a route. As highlighted in our background on game theory in Section 3, nodes usually act according to their own interests, which can be orthogonal to the overall optimal equilibrium. A game theoretic measure used by Roughgarden & Tardos in the context of routing is the *Price of Anarchy* [64], which quantifies how much a system degrades due to selfish behavior. More formally, assuming we have a measure of the efficiency of each outcome, the Price of Anarchy is the ratio between the equilibrium and the optimal outcome. Inspired by this measure, Grossklags *et al.* [65] introduce the *Price of Uncertainty*,

which measures the cost of incomplete information compared to that of complete information. An important observation is that, assuming fixed possible losses, which is reasonable in the case of mining where one can (at most) lose the fixed-cost hardware (and electricity) or stake, the more players are in the network, the less information matters. This also ties in to the value of information (*i.e.*, the possible change in utility from gaining information), which is defined for a computational setting by Halpern & Pass [66].

5.3 Consensus

5.3.1 Fault tolerance with rational players

We now look at the example of consensus. The approach here is to use incentives to bypass impossibility results on Byzantine agreement or improve on existing constructions. In order to apply GT to DS, additional adjustments have to be made. For example, traditional GT considers deviation from only one agent (as in a NE), while in practice, agents form coalition. In addition, in a DS it is important to consider multiple types of failures (*e.g.*, processors may crash) that are not considered in GT.

In order to account for both of these requirements, the BAR model defined by Aiyer *et al.* [67] introduces three different types of players: Byzantine, altruistic (players that simply follow the rules) and rational players. In this case, the expected utility of a rational player is usually defined by considering the worst configuration of Byzantine players and the worst set of strategies that those Byzantine players could take, assuming all other non-Byzantine players obey the specified strategy profile. The goal of the BAR model is to provide guarantees similar to those of Byzantine fault tolerance to all rational and altruistic nodes, as opposed to all correct nodes. Two classes of protocols meet this goal: Incentive-Compatible Byzantine Fault Tolerant (IC-BFT) protocols, and Byzantine Altruistic Rational Tolerant (BART) protocols. IC-BFT protocols, which are a subset of BART protocols, ensure that the protocol satisfies security and is the optimal one for rational nodes, while a BART protocol simply ensures security properties.

Groce *et al.* [68] introduce similar notions—perfect and statistical security—which state that in the presence of a rational adversary, the protocol still satisfies the security properties (*e.g.*, consistency and correctness for consensus). They show feasibility results of information-theoretic (both perfect and statistical) Byzantine Agreement, assuming a rational adversary and complete or partial knowledge of the adversary

preferences. Their protocols are also more efficient than traditional Byzantine Agreement protocols.

In the DAMD setting [69], participants are split into obedient, faulty, strategic, and adversarial nodes of the network. The split follows the same lines as that of the BAR model, but separates the adversarial nodes from those that are faulty with no strategic goal. Computational restrictions here are expressed with regards to the solution concepts rather than the agents. This ties into topics in computational game theory, as a solution to a DAMD problem requires not only that incentive compatibility is achieved, but also that the solution be computationally tractable, which is not always the case. (The tractability of computing Nash equilibria, or approximations, is out of the scope of this paper.) The takeaway is that many solutions on paper are not straightforwardly obtained in an algorithmic setting, whether centralized or decentralized, and even approximations may not be enough.

5.3.2 Robustness

When it comes to adapting a NE to consider coalitions and irrational players, Abraham *et al.* [70] extend the work of Halpern & Teague [34] to consider multiple players. They introduce the concept of robustness that encompasses two notions, resilience and immunity. Resilience captures the fact that a coalition of players has no incentive to deviate from the protocol, and is similar to the concept of collusion-proof NE [9]. Immunity captures the fact that even if some irrational players are present in the system, the utilities of the other players are not affected. An equilibrium that is both resilient to coalitions of up to k players, and immune to up to t irrational players, is then said to be (k, t) -robust.

Robustness is a very strong property, but it is hard to achieve in practice. Clement *et al.* [71] show that no protocol is (k, t) -robust if any node may crash, and communication is necessary and costly. When designing cryptocurrencies, however, it is not unusual to consider that communication is free.

As discussed in Section 4.1 with covert adversaries, it can be helpful to add a form of punishment to enforce correct behavior by rational players. Halpern *et al.* [70] define a (k, t) -punishment strategy such that for any coalition of at most k players and up to t irrational players, as long as more than t players use the punishment strategy and the remaining players play the equilibrium strategy, then if up to k players collude, they will be worse off than they would have been if the rational players had played the equilibrium strategy. The idea is that having more than t players use the punishment

strategy is enough to stop k players from colluding and deviating from the equilibrium strategy.

5.3.3 Price of Malice

As systems realistically involve rational and irrational players, it is important to consider how rational players react to the presence of irrational players. Moscibroda *et al.* [72] do this by considering a system with only rational and Byzantine players. They differentiate between an *oblivious* and *non oblivious model* (i.e., whether selfish players know the existence of Byzantine players or not). They define a *Byzantine Nash equilibria* that extends NE in the case where irrational players are present. In a Byzantine Nash equilibria, no selfish player can reduce their perceived expected cost (which depends on their information) by changing their strategy, given that the strategies of all other players are fixed.

In GT and MD, a concept very often discussed is the Price of Anarchy [11], which was introduced in the case of selfish routing. Moscibroda *et al.* [72] extend this to their setting, by defining the *Byzantine Price of Anarchy* that quantifies how much an optimal system degrades due to selfish behavior, when malicious players are introduced. More formally, it is the ratio between the worst social cost of a Byzantine Nash equilibrium and the minimal social cost, where the social cost of a strategy profile is the sum of all individual costs (i.e., the optimality of each outcome).

The price of Malice is used to see how a system of purely selfish players degrades in the presence of malicious irrational players. More formally, it is the ratio between the worst Byzantine Nash Equilibrium with malicious players and the Price of Anarchy in a purely selfish system.

Moscibroda *et al.* [72] also introduced the idea that Byzantine players can improve the overall system, which they called the *fear factor*. The intuition is that the rational players will adapt their strategies by fear of the actions of irrational players, rendering the overall system better. The example they introduce where this can be observed is virus inoculation. Based on the assumption that some players are irrational and will not get vaccinated, rational players will be incentivized to get vaccinated. In the case where everyone is rational, there is no equilibrium since, as long as enough people get vaccinated, the rest of the population is safe. Thus irrational players here can make the overall system better.

6. Blockchains

We now consider blockchain-based cryptocurrencies, which are an important example of systems involving aspects of both traditional security and game-theoretic aspects. In this section, we review the work that has been done by the security and distributed systems communities on blockchains that consider game-theoretic notions.

We start by reviewing the works that give some game-theoretic analysis of Bitcoin. We then illustrate how these analyses fell short with attacks that have been found on Bitcoin's incentives. We then give an overview of the work that has been done on blockchain consensus protocols that considers the question of incentives, and tries to thwart these attacks. As we argued in [Section 3](#), having a system that is secure with some bounded number of Byzantine faults is not enough to have a decentralized system as decentralization cannot be assumed. Rather, incentives should be designed to ensure enough participation and prevent coalitions. We therefore also discuss work focusing on incentivizing decentralization. We then review the work on payment channels before finally presenting some notions of fairness with respect to blockchains' reward systems.

Along the way, we also highlight new concepts of interest that have been introduced in this field, as well as how they relate to what we have previously discussed in this paper.

6.1 Game theoretic analysis of Bitcoin

Nakamoto's original Bitcoin paper [\[73\]](#) provided only informal security arguments, but several papers have since formally argued the security of Bitcoin in different models [\[74\]\[75\]\[76\]](#), usually based in the simulation setting presented in [Section 3](#) (but without a consideration of incentives).

In early work in this area, Kroll *et al.* [\[77\]](#) show that there is a NE in which all players behave consistently with Bitcoin's reference implementation, along with infinitely many equilibria in which they behave otherwise (*e.g.*, where they all agree to change a rule). Attacks like selfish mining [\[3\]\[78\]\[79\]](#) put this into question, showing that their model did not encompass behavior that could realistically occur. More recently, Fiat *et al.* [\[80\]](#) showed that the only possible pure equilibria in Bitcoin's mining are very chaotic (miners quitting and starting again periodically) or non-existent, depending on the configuration of players.

More recently Garay *et al.* [\[44\]](#) proved the security of Bitcoin in the RPD framework that was introduced in [Section 4.1](#). Their approach is based on the observation that

Bitcoin works despite its flaws, and they prove that Bitcoin is secure by relying on the rationality of players rather than an honest majority. This model inherits the flaws discussed in [Section 4.1](#) (*e.g.*, they do not consider fully malicious players). Their model also does not encompass attacks on Bitcoin's incentive structure that we now describe.

6.2 Attacks on incentives

With time, many attacks related to incentives in cryptocurrencies have been found, typically involving either external incentives (*e.g.*, in bribery attacks) or the unintended use of a cryptocurrency's technical mechanisms. The effect of these attacks results in lowering the power required for a 51% attack to less than 51%.

Selfish mining [\[3\]](#) involves a rational miner increasing their expected utility by withholding their blocks instead of broadcasting them to the rest of network, giving them an advantage in solving the new proof-of-work and making the rest of the network waste computation by mining on a block that is not the top of the chain.

Inspired by techniques introduced by Gervais *et al.* [\[81\]](#), Sapirshtein *et al.* [\[78\]](#) use Markov Decision Processes (MDP) to find the optimal strategy when doing selfish mining. (MDP are used to help make decisions in a discrete state space where outcomes are partially random.) They show that with this strategy, an adversary could mount a 51% attack with less than 25% of the computational power. This problem is further studied by Hou *et al.* [\[82\]](#) using deep reinforcement learning. They suggest that selfish mining becomes less effective when performed by multiple adversaries. In addition to withholding their own block, miners are neither incentivized to propagate information (*e.g.* transactions or blocks) to the rest of network.

This is a problem that also exists in any P2P systems [\[83\]\[84\]](#) that researchers have looked into solving, using techniques similar to those proposed in distributed systems [\[85\]\[86\]\[87\]](#).

Another issue is the verifier dilemma [\[4\]](#), which shows that miners are not incentivized to verify the content of blocks, especially when this incurs an important computation on their end.

Mining gaps are another type of attack on incentives [\[88\]\[89\]](#) where the time between the creation of blocks increases because miners wait to include enough transactions (in order to get the transaction fees).

Bribery attacks are another family of attacks that are often thought of as an example of the *tragedy of the commons*, which describes a situation where individuals acting selfishly affect the common good [90]. In our context, it captures the fact that miners have to balance their aim to maximize their profit with the risk of affecting the long term health of the cryptocurrency they mine, potentially reducing its price and their profit.

Bonneau [91] first proposed that an adversary could mount a 51% attack at a much reduced cost by renting the necessary hardware for the length of the attack rather than purchasing it. More generally, a briber could pay existing miners to mine in a certain way, without ever needing to acquire any hardware. This led to a series of papers [92][93][94][95][96] that show it is possible to introduce new incentives to an existing cryptocurrency, internally or externally, in ways that do not require trust between miners and briber (*e.g.*, using smart contracts).

Ethereum's uncle reward mechanism (that allows blocks that were mined but not appended to the blockchain to later be referenced in another block for a reward) can be used to subsidize the cost of bribery attacks [95] and selfish mining [97][98]. This is unfortunate, as they were originally introduced to aid decentralization [99] but have now been found to introduce incentives that work against this, by reducing the mining power required to perform certain attacks.

This puts into question the value of saying that a cryptocurrency is incentive-compatible if new incentives can later be added. A cryptocurrency also does not exist in a vacuum, and external incentives can always manifest in adversarial ways. Goldfinger attacks, proposed by Kroll *et al.* [77], involve an adversary paying miners of a cryptocurrency to sabotage it by mining empty blocks. In some cases, even the threat of this type of attack can be enough to kill off a cryptocurrency, as users would not want their investments to disappear if the attack happens, and thus would not invest. As a Goldfinger attack can be implemented through a smart contract in another cryptocurrency [95], it is not inconceivable that this could be attempted in practice. This clearly shows that incentives from outside the cryptocurrency itself must be considered.

Budish [100] proposes an economic analysis of 51% attacks and double spending, and shows that the Nakamoto consensus has inherent economic limitations. In particular, he shows from a strictly economic point of view that the security of the blockchain relies on scarce, non-repurposable resources (*i.e.*, ASICs) used by miners as opposed to Nakamoto's vision of "one-CPU-one-vote," and that the blockchain is vulnerable to

sabotage at a cost linear in the amount of specialized computational equipment devoted to its maintenance.

6.3 Other blockchain consensus protocols

As an alternative to existing systems (like Bitcoin and Ethereum) that have been shown to be vulnerable to the attacks we have just described, systems based on BlockDAGs rather than blockchains have been proposed [101][102][103][104][105][106]. In this model, the data structure is a Directed Acyclic Graph (DAG) of blocks, meaning that each block can have more than one parent block. When creating a new block, a miner points to all the blocks that they are aware of, revealing their view of the blockchain. This exposes more of the decision making of the players, and relates to the idea of hidden actions discussed in [Section 5](#).

Due to some additional inherent flaws in Bitcoin (*e.g.*, scalability and energy consumption), new design papers are constantly proposed by both the academic community and industry, but many leave the treatment of incentives as future work. In particular, very few papers propose an incentive scheme associated with their consensus protocols [107][108][109][105]. Moreover, the solution concepts considered in these papers are often overly-simplistic; *e.g.*, some coalition proof NE that does not consider the impact of irrational players [107][108][110][109]. Only Solidus [86] and Fantômette [105] consider robustness (introduced in [Section 5](#)).

In his draft work about incentives in Casper [111], Buterin introduces the *griefing factor* which is the ratio of the penalty incurred to the victim of an attack and the penalty incurred by the attacker. The idea of a griefing factor intuitively makes sense, as disputes in the real world can be resolved by fining a party according to the damages caused, and from a modelling point of view gives a quantifiable punishment that can be explicitly taken into account when computing equilibria. He also proves that following the protocol in Casper is a NE as long as no player holds more than a third of the deposit at stake.

Due to the lack of formal model, it can be expected that more incentive-related attacks will be proposed. For example, attacks on cryptocurrencies using PoS are now already appearing [112][113][114], further highlighting the need for better models.

In addition to the consensus rules, another route to improving the incentivization of cryptocurrencies is through their transaction fees market. As pointed out by Lavi *et al.* [115], “competition in the fee market is what keeps the rational behavior of Bitcoin’s users (partially) aligned with the goal of buying enough security for the entire

system,” and is thus crucial for its security. This problem is related to that of auction theory [116] and some of the literature of that field could be used here.

6.4 Incentivizing decentralization

Bitcoin has evolved to become different, in many ways, from the intended design and the idea of “one-CPU-one-vote” as envisioned by Nakamoto. Because the price of mining has increased exponentially with the popularity of Bitcoin, miners have started forming mining pools, where they join their resources to mine, together, more blocks.

This is obviously a big threat to the security of cryptocurrencies, as this could enable 51% attacks, which have already happened to other cryptocurrencies. As of November 2019, the most important 51% attack has targeted Ethereum Classic, which is the 16th largest cryptocurrency by market cap [117].

The centralization of cryptocurrencies has been empirically analyzed by Gencer *et al.* [118], who measured how decentralized Bitcoin’s and Ethereum’s network are. They found that three or four mining pools control more than half of the hash power of the network. This highlights the need for further research studying this occurrence of centralization and how decentralization can be maintained in practice.

Several papers propose a game-theoretic analysis of the mining pools. Arnosti *et al.* [119] model hardware investments from miners as a game, Leonardos *et al.* [120] model mining as an *oceanic game*, used to analyze decision making in settings with small numbers of big players and large numbers of individually insignificant players. Lewenberg *et al.* [121] model the mining game as a *transferable utility coalitional game*, which allows players to form coalitions and to divide their payoffs amongst themselves. As a solution concept, they use the *core*, the set of feasible allocations that cannot be improved upon by a coalition, which describes stability in coalitional games. It captures the condition under which the agents would want to form coalitions rather than not (*i.e.*, whether there exist any sub-coalition where agents could have gained more on their own). This concept is often opposed to the *Shapley value* in game theory, which defines a fair way to divide the payment among the members of a coalition based on their respective contribution, but without any consideration for stability, unlike the core. Lewenberg *et al.* additionally define the *defection* function that captures the fact that not every agent subset can collaborate and form a new coalition. They show that mining pools are generally unstable—no matter how the revenue is shared, some miners would be incentivized to switch to a different pool.

Eyal [5] also studies the stability of mining pools and proposes an attack where pools infiltrate other pools to sabotage them by joining the pool and earning rewards, but without actually contributing (*i.e.*, not revealing when they find a PoW solution). There exists configurations in which this attack constitutes both a NE and an example of a tragedy of the commons.

Mining pools can also attack each other through distributed denial of service (DDoS) attacks to lower the expected success of a competing pool (large ones in particular), rather than increasing their own computational power [122]. Over a two year period, Vasek *et al.* [123] found that 62.5% of mining pools accounting for more than 5% of the Bitcoin network power had been targeted, while only 17.1% of the smaller pools had been targeted. This has general implications for the mining ecosystem, as a peaceful equilibrium would require an increase to the cost of attacks and to the miner migration rate (miners switching pools), with no pool being significantly more attractive than others [124].

Brünjes *et al.* [125] introduce and study reward sharing schemes that promote the fair formation of stake pools in a PoS blockchain. They argue that a NE only considers *myopic players*, *i.e.*, players who ignore the responses to their own actions. As a result, they consider the notion of *non-myopic Nash equilibrium* (based on previous work by Fiat *et al.* [126]), which captures the effects that a certain move will incur, anticipating a strategic response from the other players.

Luu *et al.* [127] use smart contracts to decentralize mining by incurring mining fees lower than centralized mining pools. Miller *et al.* [128] present several definitions and constructions for “non-outsourcable” puzzles. Both papers use informal arguments to justify their construction as opposed to a formal model.

In a recent paper, Kwon *et al.* [129] propose a formal model for the decentralization of blockchains, and show that full decentralization is impossible unless there exists a Sybil cost.

6.5 Payment Channels

In order to overcome the scalability issues of Bitcoin, a new concept, referred to as *layer 2* or *payment channels* has been proposed [130]. The idea is that, since the network cannot handle enough transactions, participants can take some transactions off-chain (*i.e.*, outside the main network) by opening a channel between themselves. This is done by locking a deposit on the blockchain, opening the channel and transacting on the channel, then settling the overall balance of all transactions on-

chain so that the blockchain will see only two transactions (locking the funds and settling the balance).

Several designs have been proposed to achieve this [130][131][132]. The high-level idea is that participants will create evidence of each of their transactions (*e.g.*, using signatures), so that whenever someone tries to cheat, the other party can prove it and receive the cheating party's deposit as compensation.

In this setting, the security relies on the fact that cheating is easily detectable due to cryptographic evidence, and on the financial punishment associated with it. So again, incentives are tightly linked to security. A few papers [131][133] present formal models to analyze the security of these payment channels. They are based on the UC model mentioned in [Section 3](#) but do not consider utilities (although it is an important part of the security of the system).

In order to facilitate payment channels, a routing solution has been proposed [134]. There are usually difficulties in this case, due to the need for collaterals to be locked by everyone on the routing path. This work is related to the one on selfish routing discussed in [Section 5](#).

A problem with payment channels is the requirement for participants to be online to detect cheating (*i.e.*, the cheater broadcasting an old balance to the blockchain). McCorry *et al.* [135] propose delegating this task to a third party, a *watchtower*, but it is unclear how incentives should be designed in this context.

6.6 Fairness

Fairness in cryptocurrencies is implicitly captured by the notion of chain quality introduced by Garay *et al.* [76], which states that an adversary should not contribute more blocks to the blockchain than what they are supposed to (*i.e.*, proportionally to their computational power in the PoW setting).

Chen *et al.* [136] recently showed that a proportional reward system is the unique allocation rule that satisfies properties of symmetry, budget balance (weak or strong), sybil-proofness, and collusion-proofness, which are desirable.

In the PoS setting, Fanti *et al.* [112] define a notion of *equitability* that corresponds to how much a node's initial investment (*i.e.*, stake) can grow or shrink, and address the problem of the "rich get richer" in PoS cryptocurrencies (which arguably also exists in PoW cryptocurrencies). They propose a geometric reward function that they prove is

more equitable (*i.e.*, the distribution of stake stays more stable). In general, the problem of compounding of wealth is reinforced by the fact that early adopters of a cryptocurrency have a significant advantage, benefiting from the ease of mining (or staking) and greatly cheaper coin prices in the early days. Dealing with this is more of a macroeconomic problem that, to the best of our knowledge, has not yet received any attention.

7. Discussion

Within each section of this paper, we have reviewed models that draw on game theoretic tools and could thus be used to address problems in blockchain based cryptocurrencies. [Table 1](#) gives a summary of some these models, and we will now discuss general points and lessons that can be drawn from the work so far.

Concept	Player types	Description	Shortcomings	Example, suggestions, and discussion
RPD	Rational, honest	Meta-game between the game designer and the adversary, which the adversary wins if it can exploit a vulnerabilities without decreasing its utility.	Does not consider irrational adversaries.	Could be used in Layer 2 since in that case a misbehaving participants will be penalized and the other participant will earn all of its money so Byzantine adversaries cannot harm honest players more than rational adversaries.
BMG	Rational	Considers games and the complexity of actions by modelling players as Turing machines	Results are primarily equivalence theorems with dense proofs rather than new results or simpler proof methods. No consideration of Byzantine adversaries.	This could be usefully applied to cryptocurrencies given that many situations are modeled as games involving computation e.g., mining.
Price of Unaccountability	Rational	Worst ratio between utilities in the observable-actions case and the hidden actions case.	Defined in the principal-agent model.	As differentiating between malicious behavior and genuine latency is hard, especially in PoS systems, the Price of Unaccountability may be a useful to evaluating this.
BAR	Byzantine, rational, honest	Introduces the idea of BART and IC-BFT protocols, which make it possible to bypass impossibility results in consensus protocols.		
(k, t) -robustness	Byzantine, rational, honest	Extension of NE that consider coalitions of both rational and Byzantine players.	The two concepts of immunity and robustness are treated separately.	This model is better suited than NE to study the game theoretic aspects of consensus protocols that involve many players and coalitions.
Price of Malice and Byzantine Price of Anarchy	Byzantine, rational	Quantifies how much a system degrades with the presence of irrational players and relates to the concept of immunity introduced in (k, t) -robustness.	Does not capture how different information sets impact the system.	Could inspire new measures that quantify the trade-off between blockchain and traditional consensus protocols. Blockchain-based systems are intended to be more scalable as they are meant to handle open participation, compared to classical consensus that requires the many messages to be exchanged, but in the case of Bitcoin this comes at the price of PoW so there is an incurred economic cost.
Fear factor	Byzantine, rational	Rational players are incentivized to follow the protocol by fear of Byzantine players. Similar to Price of Malice (if the system improves instead of deteriorates).	Same as above.	This could be used to argue against the verifier dilemma. Some users may be motivated to verify the content of blocks by fear that others will not.

Image 1
Table 1

The most immediate problem that most models try to address is the consideration of incentives at the level of security properties by encoding utility functions and some notion of equilibrium into a modified definition of the traditional security property. For example, RPD does this for UC security, and the BAR model does this for distributed systems. Already, however, some differences emerge.

First, there are differences in the types of players. Different models assume players can be some subset of honest, rational, or Byzantine. How well the assumed types of

players reflect the reality of the system will have a great impact on the usefulness of the model, regardless of its technical merits.

In particular, while honest and Byzantine players can be defined with respect to a protocol and whether or not they follow it, the meaning of a rational player is harder to pin down. This is because a player is, in practice, rational not only with respect to a protocol and the incentives in the system that the protocol is part of, but also with respect to a variety of possible external incentives. We have referred to this with respect to rational cryptography and bribery attacks, but it is a more general point that Ford & Böhme have explicitly highlighted [\[137\]](#).

Security models are built such that within the model, properties can remain true up to some amount of adversarial capabilities (*e.g.*, a third of Byzantine nodes and computational capabilities), but it is not clear how valid the models are when the assumed player types differ from reality. Moreover, comparing the models presented in this paper would require a concrete idea of what the correct assumptions that can be made about participants in cryptocurrencies are. Each model is constructed to work better than others within the constraints of their assumptions. As remarked by Box, “all models are wrong, but some are useful.” [\[138\]](#) It seems that the current limitation of the existing literature is in understanding how useful these models are in practice rather than in introducing new models.

In Economics, Becker pointed out long ago that it was not always clear what rationality implied (because some observed behavior was compatible with both rational and irrational behavior [\[139\]](#)), and concluded that perhaps irrationality deserved to be studied with more attention. Behavioral Economics and, more broadly, experimental data-driven work has taken an important place in Economics to understand why certain models did not work in practice. Perhaps the same should be done here.

For example, mining rewards can be designed as part of a protocol to ensure some level of decentralization, but the utility function of miners will also depend on their individual economic environment that the protocol cannot fix. In this case, relating the economics of miners to their relationship with the system could do more to ensure some level of decentralization than tweaking the rewards given out by the protocol.

Second, there are the way coalitions of players are treated. Pools are being observed and studied carefully in the community, but some important problems remain understudied. It could be argued that, in some cases, pools somehow self-regulate as in, for example, the case of the Ghash.io Bitcoin pool that once got more than 51% of the

hashing power but then decided to withdraw part of it [\[140\]](#). This could be because of the fear of a loss of confidence by users of the system, which, if it is justified, would mean that the incentives were (to some extent) aligned to prevent a 51% attack.

Coalitions are also easily abstracted as an entity controlling a fixed share of power in the system, but this ignores costs inherent to colluding such as, for example, communications costs as a coalition must reach some form of consensus on what actions it takes. In the same way that results are obtained for a whole system based on assumed proportions of player types in the system, perhaps useful results can be obtained based on the proportions of player types in a coalition with respect to that coalition.

8. Conclusion

Security researchers and cryptographers have been interested in incorporating game-theoretic notions to their models for many years. In this work, we have highlighted existing concepts and explained how and where they could be used for specific applications.

The approach taken in most of the papers that we described here is to extend a field by, for example, incorporating utility functions (rational cryptography) or computation (Bayesian machine games). No completely new theory has appeared, and it would be interesting to see a new theory built from the ground up to address considerations of incentives at all stages of the design process, rather than adapting existing models. We hope that this paper will give some inspiration towards new formal models.

Acknowledgments

The authors would like to thank their anonymous reviewers for their useful feedback. Alexander Hicks is supported by OneSpan¹ and UCL through an EPSRC Research Studentship.

Appendix

A. Appendix

In this appendix, we provide formal definitions for some of the concepts presented in the main body of the paper that are not formally defined.

A.1 Game Theory

To start off, we introduce the standard definitions for Bayesian games and mechanisms.

- **Bayesian game setting:** A Bayesian game setting is a tuple $(N, O, \Theta, \text{Pr}, u)$, where:
 - N is a finite set of n players;
 - O is a set of outcomes;
 - $\Theta = \Theta_1, \dots, \Theta_n$ is a set of possible joint type vectors;
 - Pr is a (common prior) probability distribution on Θ ; and
 - $u = (u_1, \dots, u_n)$, where $u_i : O \times \mathbb{R} \rightarrow \mathbb{R}$ is the utility function for each player i .
- **Mechanism for a Bayesian game setting:** A mechanism for a Bayesian game setting (N, O, Θ, p, u) is a pair (A, M) , where:
 - $A = A_1 \times \dots \times A_n$, where A_i is the set of actions available to agent $i \in N$; and
 - $M : A \rightarrow \mathcal{D}(O)$ maps each action profile to a distribution over outcomes

A.2 Game Theory and Cryptography

We now move on to concepts presented in [Section 4](#).

- **ϵ -subgame perfect equilibrium [43]:** Let $\mathcal{G}_{\mathcal{M}}$ be an attack game. A strategy profile (A, Π) is an ϵ -subgame perfect equilibrium in $\mathcal{G}_{\mathcal{M}}$ if: (1) for any $\Pi' \in \text{ITM}^n$, $u_D(\Pi', A(\Pi')) \leq u_D(\Pi, A(\Pi)) + \epsilon$; and (2) for any $A' \in \text{ITM}$, $u_A(\Pi, A'(\Pi)) \leq u_A(\Pi, A(\Pi)) + \epsilon$.
- **Attack-payoff security [43]:** Let $\mathcal{M} = (\mathcal{F}, \langle \mathcal{F} \rangle, v)$ be an attack model and let Π be a protocol that realizes functionality $\langle \mathcal{F} \rangle$. Π is attack-payoff secure in $\mathcal{M}$ if $\vec{U}^{\Pi, \langle \mathcal{F} \rangle} \stackrel{\text{negl}}{\leq} \vec{U}^{\Phi^{\mathcal{F}}, \langle \mathcal{F} \rangle}$ where $\Phi^{\mathcal{F}}$ is the “dummy” $\mathcal{F}$ hybrid protocol (*i.e.*, the protocol that forwards all inputs and outputs from the functionality $\mathcal{F}$, see [Section 3](#)) and $\vec{U}^{\Pi, \langle \mathcal{F} \rangle}$ is the maximized ideal expected payoff of an adversary.
- **Incentive compatibility [44]:** Let Π be a protocol and Pr be a set of PT protocols that have access to the same hybrids as Π . We say that Π is Pr -incentive compatible in the attack model $\mathcal{M}$ if and only if for some $\text{Adv}(\Pi, \text{Adv})$ is a (Pr, ITM) -subgame perfect equilibrium in the attack game defined by $\mathcal{M}$.
- **Bayesian Machine Game [46]:** A Bayesian machine game G is described by a tuple $(N, \mathcal{M}, \Theta, \text{Pr}, \mathcal{C}_1, \dots, \mathcal{C}_m, u_1, \dots, u_2)$ where:
 - N is the set of players, $\mathcal{M}$ is the set of possible machines;
 - $\Theta \subseteq (\{0, 1\}^*)^{m+1}$ is the set of type profiles where the $(m+1)$ st element in the profile corresponds to nature’s type;
 - Pr is a distribution on Θ ;
 - $\mathcal{C}_i$ is a complexity function; and
 - $u_i : T \times (\{0, 1\}^*)^m \times \mathbb{N} \rightarrow \mathbb{R}$ is player i ’s utility function.

Given a Bayesian machine game G , a machine profile $\vec{M}$, and $\epsilon \geq 0$, M_i is an ϵ -best response to $\vec{M}_{-i}$ (the tuple consisting of all machines in $\vec{M}$ other than M_i) if, for every $M'_i \in \mathcal{M}$,

$$U_i^G[(M_i, \vec{M}_{-i})] \geq U_i^G[(M'_i, \vec{M}_{-i})] - \epsilon. \quad (1)$$

$\vec{M}$ is an ϵ -Nash equilibrium of G if, for all players i , M_i is an ϵ -best response to $\vec{M}_{-i}$. A Nash equilibrium is a 0-Nash equilibrium.

- **Universal implementation [46]:** Suppose that $\mathcal{G}$ is a set of n -player canonical games, $\mathcal{Z}$ is a subsets of N , $\mathcal{F}$ and $\mathcal{F}'$ are mediators, $M_1, \dots, M_n$ are interactive machines, $p : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$ and $\epsilon : \mathbb{N} \rightarrow \mathbb{R}$. $(M, \mathcal{F}')$ is a $(\mathcal{G}, \mathcal{Z}, p)$ -universal implementation of $\mathcal{F}$ with error ϵ if, for all n , all games $G \in \mathcal{G}$ with input length n and all $\mathcal{Z}' \subseteq \mathcal{Z}$ if $\vec{\Lambda}^{\mathcal{F}}$ is a $p(n, \cdot)$ -robust $\mathcal{Z}'$ -safe ϵ -NE in the mediated machine game $(G, \mathcal{F})$ then
 1. (Preserving equilibrium) $\vec{M}$ is a $\mathcal{Z}'$ -safe ϵ -NE in the mediated machine game $(G, \mathcal{F}')$
 2. (Preserving Action Distributions) For each type profile $\vec{t}$, the action profile induced by $\vec{\Lambda}^{\mathcal{F}}$ in $(G, \mathcal{F})$ is identically distributed to then action profile induced by M in $(G, \mathcal{F}')$.
- **Sequential equilibrium in computational games [141]:** A pair $(\vec{M}, \mu)$ consisting of a machine profile $\vec{M}$ and a belief system μ is called a belief assessment. A belief assessment $(\vec{M}, \mu)$ is an interim (resp. ex ante) sequential equilibrium in a machine game G if 1) μ is compatible with $\vec{M}$; and 2) for all players i and states q of M_i , as well as machines M'_i that are compatible with M_i and q such that $(M_i, q, M'_i) \in \mathcal{M}$ (the set of possible machines) (resp. (M_i, q, M'_i) is a local variant of M_i), we have

$$U_i(\vec{M}|q, \mu) \geq U_i(((M_i, q, M'_i), \vec{M}_{-i})|q, \mu) \quad (2)$$

A.3 Game Theory and Distributed Design

Finally, we give definitions for concepts presented in [Section 5](#).

- **Incentive-Compatible Byzantine Fault Tolerant (IC-BFT) protocols [67]:** A protocol is IC-BFT if it guarantees the specified set of safety and liveness properties and if it is in the best interest of all rational nodes to follow the protocol exactly.
- **Byzantine Altruistic Rational Tolerant (BART) protocols [67]:** A protocol is BART if it guarantees the specified set of safety and liveness properties in the presence of all rational deviations from the protocol.

- **Perfect security [68]** : A protocol for broadcast or consensus is perfectly secure against rational adversaries controlling t players with utility U if for every t -adversary there is a strategy S such that for any choice of input for honest players (1) (S is tolerable): S induces a distribution of final outputs D in which no security condition is violated with nonzero probability, and (2) (S is Nash): For any strategy $S' \neq S$ with induced output distribution $D' : U(D) \geq U(D')$.
- **Statistical Security [68]**: A protocol for broadcast or consensus is statistically secure against rational adversaries controlling t players with utility U if for every t -adversary there is a strategy S such that, for any choice of input for honest players, S induces a distribution of final outputs D_k when the security parameter is k and the following properties hold: (1) (S is tolerable): no security condition is violated with nonzero probability in D_k for any k , and (2) (S is statistical Nash): for any strategy $S' \neq S$ with induced output distributions D'_k there is a negligible function $\text{negl}(\cdot)$ such that $U(D_k) + \text{negl}(k) > U(D'_k)$.
- **(k, t)-robustness [70]**: A strategy profile σ is a (k, t) -robust equilibrium if for all $C, T \subseteq N$, $C \cap T = \emptyset$, $|C| \leq k$, $|T| \leq t \forall \tau_t \in \mathcal{S}_T \forall \phi_C \in C$ we have:
 $u_i(\sigma_{-T}, \tau_T) \geq u_i(\sigma_{-C \cap T}, \phi_C, \tau_T)$.
- **(k, t)-punishment [70]**: A joint strategy ρ is a (k, t) -punishment strategy with respect to σ if for all $C, T, P \subseteq N$ such that C, T, P are disjoint, $|C| \leq k$, $|T| \leq t$, and $|P| > t$, for all $\tau_T \in \mathcal{S}_T$, for all $\phi_C \in \mathcal{S}_C$, for all $i \in C$ we have
 $u_i(\sigma_T, \tau_T) > u_i(\sigma_{N-(C \cup T \cup P)}, \phi_C, \tau_T, \rho_P)$.

Footnotes

1. <https://www.onespan.com> [↗](#)

Citations

1. Anderson, R. (2001). Why information security is hard-an economic perspective. In *Computer security applications conference, 2001. Acsac 2001. Proceedings 17th annual* (pp. 358-365). IEEE. [↗](#)
2. Anderson, R., & Moore, T. (2006). The economics of information security. *Science*, 314(5799), 610-613. [↗](#)
3. Eyal, I., & Sirer, E. G. (2018). Majority is not enough: Bitcoin mining is vulnerable. *Commun. ACM*, 61(7), 95-102. <https://doi.org/10.1145/3212998> [↗](#)

4. Luu, L., Teutsch, J., Kulkarni, R., & Saxena, P. (2015). Demystifying incentives in the consensus computer. In *Proceedings of the 22Nd acm sigsac conference on computer and communications security* (pp. 706–719). New York, NY, USA: ACM. <https://doi.org/10.1145/2810103.2813659> [↵](#)
5. Eyal, I. (2015). The miner’s dilemma. In *Proceedings of the 2015 ieee symposium on security and privacy* (pp. 89–103). Washington, DC, USA: IEEE Computer Society. <https://doi.org/10.1109/SP.2015.13> [↵](#)
6. Linial, N. (2002). Games computers play: Game-theoretic aspects of computing. [↵](#)
7. Katz, J. (2008). Bridging game theory and cryptography: Recent results and future directions. In *Proceedings of the 5th conference on theory of cryptography* (pp. 251–272). Berlin, Heidelberg: Springer-Verlag. Retrieved from <http://dl.acm.org/citation.cfm?id=1802614.1802635> [↵](#)
8. Halpern, J. Y. (2007). Computer science and game theory: A brief survey. *CoRR, abs/cs/0703148*. Retrieved from <http://arxiv.org/abs/cs/0703148> [↵](#)
9. Halpern, J. Y. (2008). Beyond Nash equilibrium: Solution concepts for the 21st century. *CoRR, abs/0806.2139*. Retrieved from <http://arxiv.org/abs/0806.2139> [↵](#)
10. Shoham, Y. (2008). Computer science and game theory. *Communications of the ACM, 51*(8), 74–79. [↵](#)
11. Roughgarden, T. (2010). Algorithmic game theory. *Communications of the ACM, 53*(7), 78–86. [↵](#)
12. Nisan, N., Roughgarden, T., Tardos, E., & Vazirani, V. V. (2007). *Algorithmic game theory*. Cambridge University Press. [↵](#)
13. Tambe, M. (2011). *Security and game theory: Algorithms, deployed systems, lessons learned*. Cambridge University Press. [↵](#)
14. Grossklags, J., Christin, N., & Chuang, J. (2008). Secure or insure?: A game-theoretic analysis of information security games. In *Proceedings of the 17th international conference on world wide web* (pp. 209–218). New York, NY, USA: ACM. <https://doi.org/10.1145/1367497.1367526> [↵](#)

15. Bano, S., Sonnino, A., Al-Bassam, M., Azouvi, S., McCorry, P., Meiklejohn, S., & Danezis, G. (2017). Consensus in the age of blockchains. *ArXiv Preprint ArXiv:1711.03936*. [↵](#)
16. Garay, J., & Kiayias, A. (n.d.). SoK: A consensus taxonomy in the blockchain era. [↵](#)
17. Wang, W., Hoang, D. T., Xiong, Z., Niyato, D., Wang, P., Hu, P., & Wen, Y. (2018). A survey on consensus mechanisms and mining management in blockchain networks. *ArXiv Preprint ArXiv:1805.02707*. [↵](#)
18. Stifter, N., Judmayer, A., Schindler, P., Zamyatin, A., & Weippl, E. (n.d.). Agreement with satoshi-on the formalization of nakamoto consensus. [↵](#)
19. Bonneau, J., Miller, A., Clark, J., Narayanan, A., Kroll, J. A., & Felten, E. W. (2015). SoK: Research perspectives and challenges for bitcoin and cryptocurrencies. In *2015 IEEE Symposium on Security and Privacy* (pp. 104–121). <https://doi.org/10.1109/SP.2015.14> [↵](#)
20. Cachin, C., & Vukolić, M. (2017). Blockchains consensus protocols in the wild. *ArXiv Preprint ArXiv:1707.01873*. [↵](#)
21. Liu, Z., Luong, N. C., Wang, W., Niyato, D., Wang, P., Liang, Y.-C., & Kim, D. I. (2019). A survey on blockchain: A game theoretical perspective. *IEEE Access*, 7, 47615–47643. [↵](#)
22. Owen, G. (1968). *Game theory*. Saunders. Retrieved from <https://books.google.co.uk/books?id=v6lMAAAAMAAJ> [↵](#)
23. Jackson, M. O. (2003). *Mechanism theory*. California Institute of Technology. [↵](#)
24. Gneezy, U., & Rustichini, A. (2000). A fine is a price. *The Journal of Legal Studies*, 29(1), 1–17. [↵](#)
25. Goldreich, O., Micali, S., & Wigderson, A. (1987). How to play any mental game. In *Proceedings of the nineteenth annual ACM symposium on theory of computing* (pp. 218–229). New York, NY, USA: ACM. <https://doi.org/10.1145/28395.28420> [↵](#)
26. Micali, S., & Rogaway, P. (1991). Secure computation. In *Annual international cryptology conference* (pp. 392–404). Springer. [↵](#)

27. Lindell, Y. (2016). *How to simulate it - A tutorial on the simulation proof technique*. Cryptology ePrint Archive, Report 2016/046. [↵](#)
28. Canetti, R. (2001). Universally composable security: A new paradigm for cryptographic protocols. In focs01ed (Ed.), *Focs01name* (pp. 136–145). focs01addr: focspub. [↵](#)
29. Troncoso, C., Isaakidis, M., Danezis, G., & Halpin, H. (2017). Systematizing decentralization and privacy: Lessons from 15 years of research and deployments. *Proceedings on Privacy Enhancing Technologies*, 2017(4), 404–426. Retrieved from <https://content.sciendo.com/view/journals/popets/2017/4/article-p404.xml> [↵](#)
30. Azouvi, S., Hicks, A., & Murdoch, S. J. (2018). Incentives in security protocols. In V. Matyáš, P. Švenda, F. Stajano, B. Christianson, & J. Anderson (Eds.), *Security protocols xxvi* (pp. 132–141). Cham: Springer International Publishing. [↵](#)
31. *Crypto51*. (n.d.). Retrieved from <https://www.crypto51.app/> [↵](#)
32. *Ethereum*. (n.d.). Retrieved from <https://www.ethereum.org> [↵](#)
33. Dodis, Y., Halevi, S., & Rabin, T. (2000). A cryptographic solution to a game theoretic problem. In M. Bellare (Ed.), *Advances in cryptology — crypto 2000* (pp. 112–130). Berlin, Heidelberg: Springer Berlin Heidelberg. [↵](#)
34. Halpern, J., & Teague, V. (2004). Rational secret sharing and multiparty computation: Extended abstract. In *Proceedings of the thirty-sixth annual acm symposium on theory of computing* (pp. 623–632). New York, NY, USA: ACM. <https://doi.org/10.1145/1007352.1007447> [↵](#)
35. Kol, G., & Naor, M. (2008). Games for exchanging information. In *Proceedings of the fortieth annual acm symposium on theory of computing* (pp. 423–432). ACM. [↵](#)
36. Kol, G., & Naor, M. (2008). Cryptography and game theory: Designing protocols for exchanging information. In *Theory of cryptography conference* (pp. 320–339). Springer. [↵](#)
37. Gordon, S. D., & Katz, J. (2006). Rational secret sharing, revisited. In *International conference on security and cryptography for networks* (pp. 229–241). Springer. [↵](#)

38. Fuchsbauer, G., Katz, J., & Naccache, D. (2010). Efficient rational secret sharing in standard communication networks. In *Theory of cryptography conference* (pp. 419–436). Springer. [↵](#)
39. Asharov, G., Canetti, R., & Hazay, C. (2011). Towards a game theoretic view of secure computation. In K. G. Paterson (Ed.), *Advances in cryptology - eurocrypt 2011* (pp. 426–445). Berlin, Heidelberg: Springer Berlin Heidelberg. [↵](#)
40. Dodis, Y., Rabin, T., & others. (2007). Cryptography and game theory. *Algorithmic Game Theory*, 181–207. [↵](#)
41. Aumann, Y., & Lindell, Y. (2007). Security against covert adversaries: Efficient protocols for realistic adversaries. In *Theory of cryptography conference* (pp. 137–156). Springer. [↵](#)
42. Lysyanskaya, A., & Triandopoulos, N. (2006). Rationality and adversarial behavior in multi-party computation. In C. Dwork (Ed.), *Advances in cryptology - crypto 2006* (pp. 180–197). Berlin, Heidelberg: Springer Berlin Heidelberg. [↵](#)
43. Garay, J., Katz, J., Maurer, U., Tackmann, B., & Zikas, V. (2013). *Rational protocol design: Cryptography against incentive-driven adversaries*. Cryptology ePrint Archive, Report 2013/496. [↵](#)
44. Badertscher, C., Garay, J., Maurer, U., Tschudi, D., & Zikas, V. (2018). *But why does it work? A rational protocol design treatment of bitcoin*. Cryptology ePrint Archive, Report 2018/138. [↵](#)
45. Garay, J., Katz, J., Tackmann, B., & Zikas, V. (2015). *How fair is your protocol? A utility-based approach to protocol optimality*. Cryptology ePrint Archive, Report 2015/187. [↵](#)
46. Halpern, J. Y., & Pass, R. (2008). Game Theory with Costly Computation. *ArXiv E-Prints*. Retrieved from <http://arxiv.org/abs/0809.0024> [↵](#)
47. Pass, R., & Halpern, J. (2010). Game theory with costly computation: Formulation and application to protocol security. In *Proceedings of the behavioral and quantitative game theory: Conference on future directions* (p. 89:1-89:1). New York, NY, USA: ACM. <https://doi.org/10.1145/1807406.1807495> [↵](#)
48. Halpern, J. Y., & Pass, R. (2011). Algorithmic rationality: Adding cost of computation to game theory. *SIGecom Exch.*, 10(2), 9–15.

<https://doi.org/10.1145/1998549.1998551> [↵](#)

49. Halpern, J. Y., & Pass, R. (2015). Algorithmic rationality: Game theory with costly computation. *Journal of Economic Theory*, 156, 246–268. [↵](#)
50. Halpern, J. Y., Pass, R., & Reichman, D. (2015). On the Non-Existence of Nash Equilibrium in Games with Resource-Bounded Players. *ArXiv E-Prints*. Retrieved from <http://arxiv.org/abs/1507.01501> [↵](#)
51. Ben-Or, M., Goldwasser, S., & Wigderson, A. (1988). Completeness theorems for non-cryptographic fault-tolerant distributed computation. In *Proceedings of the twentieth annual acm symposium on theory of computing* (pp. 1–10). New York, NY, USA: ACM. <https://doi.org/10.1145/62212.62213> [↵](#)
52. Nisan, N., & Ronen, A. (2001). Algorithmic mechanism design. *Games and Economic Behavior*, 35(1-2), 166–196. [↵](#)
53. Feigenbaum, J., Papadimitriou, C. H., & Shenker, S. (2001). Sharing the cost of multicast transmissions. *Journal of Computer and System Sciences*, 63(1), 21–41. [↵](#)
54. Feigenbaum, J., & Shenker, S. (2002). Distributed algorithmic mechanism design: Recent results and future directions. [↵](#)
55. Varian, H. (2004). System reliability and free riding. In *Economics of information security* (pp. 1–15). Springer. [↵](#)
56. Feldman, M., & Chuang, J. (2005). Overcoming free-riding behavior in peer-to-peer systems. *ACM Sigecom Exchanges*, 5(4), 41–50. [↵](#)
57. Babaioff, M., Chuang, J., & Feldman, M. (2007). Incentives in peer-to-peer systems. *Algorithmic Game Theory*, 593–611. [↵](#)
58. Cohen, B. (2003). Incentives build robustness in bittorrent. In *Workshop on economics of peer-to-peer systems* (Vol. 6, pp. 68–72). [↵](#)
59. Piatek, M., Isdal, T., Anderson, T., Krishnamurthy, A., & Venkataramani, A. (2007). Do incentives build robustness in bittorrent. In *Proc. Of nsdi* (Vol. 7). [↵](#)
60. Jun, S., & Ahamad, M. (2005). Incentives in bittorrent induce free riding. In *Proceedings of the 2005 acm sigcomm workshop on economics of peer-to-peer systems* (pp. 116–121). ACM. [↵](#)

61. Feldman, M., Papadimitriou, C., Chuang, J., & Stoica, I. (2006). Free-riding and whitewashing in peer-to-peer systems. *IEEE Journal on Selected Areas in Communications*, 24(5), 1010–1019. [↵](#)
62. Chase, M., & Meiklejohn, S. (2016). Transparency overlays and applications. In *Proceedings of the 2016 acm sigsac conference on computer and communications security* (pp. 168–179). New York, NY, USA: ACM. <https://doi.org/10.1145/2976749.2978404> [↵](#)
63. Vishnumurthy, V., Chandrakumar, S., & Sirer, E. G. (2003). Karma: A secure economic framework for peer-to-peer resource sharing. In *Workshop on Economics of Peer-to-peer Systems* (Vol. 35). [↵](#)
64. Roughgarden, T., & Tardos, É. (2002). How bad is selfish routing? *Journal of the ACM (JACM)*, 49(2), 236–259. [↵](#)
65. Grossklags, J., Johnson, B., & Christin, N. (2010). The price of uncertainty in security games. In *Economics of information security and privacy* (pp. 9–32). Springer. [↵](#)
66. Halpern, J. Y. (2010). I don’t want to think about it now: Decision theory with costly computation. In *Twelfth international conference on the principles of knowledge representation and reasoning*. [↵](#)
67. Aiyer, A. S., Alvisi, L., Clement, A., Dahlin, M., Martin, J.-P., & Porth, C. (2005). BAR Fault Tolerance for cooperative services. *SIGOPS Oper. Syst. Rev.*, 39(5), 45–58. <https://doi.org/10.1145/1095809.1095816> [↵](#)
68. Groce, A., Katz, J., Thiruvengadam, A., & Zikas, V. (2012). Byzantine agreement with a rational adversary. In A. Czumaj, K. Mehlhorn, A. Pitts, & R. Wattenhofer (Eds.), *Automata, languages, and programming* (pp. 561–572). Berlin, Heidelberg: Springer Berlin Heidelberg. [↵](#)
69. Feigenbaum, J., Schapira, M., & Shenker, S. (2007). Distributed algorithmic mechanism design. *Algorithmic Game Theory*, 363–384. [↵](#)
70. Abraham, I., Dolev, D., Gonen, R., & Halpern, J. (2006). Distributed computing meets game theory: Robust mechanisms for rational secret sharing and multiparty computation. In *Proceedings of the twenty-fifth annual acm symposium on principles of distributed computing* (pp. 53–62). New York, NY, USA: ACM. <https://doi.org/10.1145/1146381.1146393> [↵](#)

71. Clement, A., Napper, J., Li, H., Martin, J., Alvisi, L., & Dahlin, M. (2007). Theory of BAR games. In *Brief announcements: Proceedings of the symposium on principles of distributed computing (podc 2007)*. [↵](#)
72. Moscibroda, T., Schmid, S., & Wattenhofer, R. (2006). When selfish meets evil: Byzantine players in a virus inoculation game. In *Proceedings of the twenty-fifth annual acm symposium on principles of distributed computing* (pp. 35–44). New York, NY, USA: ACM. <https://doi.org/10.1145/1146381.1146391> [↵](#)
73. Nakamoto, S. (2008). *Bitcoin: A peer-to-peer electronic cash system*. [↵](#)
74. Pass, R., Seeman, L., & Shelat, A. (2017). Analysis of the blockchain protocol in asynchronous networks. In *Annual international conference on the theory and applications of cryptographic techniques* (pp. 643–673). Springer. [↵](#)
75. Kiffer, L., Rajaraman, R., & shelat, abhi. (2018). A better method to analyze blockchain consistency. In *Proceedings of the 2018 acm sigsac conference on computer and communications security* (pp. 729–744). New York, NY, USA: ACM. <https://doi.org/10.1145/3243734.3243814> [↵](#)
76. Garay, J., Kiayias, A., & Leonardos, N. (2015). The bitcoin backbone protocol: Analysis and applications. In E. Oswald & M. Fischlin (Eds.), *Advances in cryptology - eurocrypt 2015* (pp. 281–310). Berlin, Heidelberg: Springer Berlin Heidelberg. [↵](#)
77. Kroll, J. A., Davey, I. C., & Felten, E. W. (2013). The economics of bitcoin mining , or bitcoin in the presence of adversaries. [↵](#)
78. Sapirshtein, A., Sompolinsky, Y., & Zohar, A. (2017). Optimal selfish mining strategies in bitcoin. In J. Grossklags & B. Preneel (Eds.), *Financial cryptography and data security* (pp. 515–532). Berlin, Heidelberg: Springer Berlin Heidelberg. [↵](#)
79. Nayak, K., Kumar, S., Miller, A., & Shi, E. (2016). Stubborn mining: Generalizing selfish mining and combining with an eclipse attack. In *Security and privacy (euros&P), 2016 ieee european symposium on* (pp. 305–320). IEEE. [↵](#)
80. Fiat, A., Karlin, A., Koutsoupias, E., & Papadimitriou, C. (2019). Energy equilibria in proof-of-work mining. In *Proceedings of the 2019 acm conference on economics and computation* (pp. 489–502). New York, NY, USA: ACM. <https://doi.org/10.1145/3328526.3329630> [↵](#)

81. Gervais, A., Karame, G. O., Wüst, K., Glykantzis, V., Ritzdorf, H., & Capkun, S. (2016). On the security and performance of proof of work blockchains. In *Proceedings of the 2016 acm sigsac conference on computer and communications security* (pp. 3–16). New York, NY, USA: ACM. <https://doi.org/10.1145/2976749.2978341> [↵](#)
82. Hou, C., Zhou, M., Ji, Y., Daian, P., Tramer, F., Fanti, G., & Juels, A. (2019). SquirRL: Automating attack discovery on blockchain incentive mechanisms with deep reinforcement learning. *ArXiv Preprint ArXiv:1912.01798*. [↵](#)
83. Emek, Y., Karidi, R., Tennenholtz, M., & Zohar, A. (2011). Mechanisms for multi-level marketing. In *Proceedings of the 12th acm conference on electronic commerce* (pp. 209–218). ACM. [↵](#)
84. Li, C., Yu, B., & Sycara, K. (2009). An incentive mechanism for message relaying in unstructured peer-to-peer systems. *Electronic Commerce Research and Applications*, 8(6), 315–326. [↵](#)
85. Babaioff, M., Dobzinski, S., Oren, S., & Zohar, A. (2012). On bitcoin and red balloons. In *Proceedings of the 13th acm conference on electronic commerce* (pp. 56–73). ACM. [↵](#)
86. Abraham, I., Malkhi, D., Nayak, K., Ren, L., & Spiegelman, A. (2016). Solidus: An incentive-compatible cryptocurrency based on permissionless Byzantine consensus. *CoRR*, abs/1612.02916. Retrieved from <http://arxiv.org/abs/1612.02916> [↵](#)
87. Ersoy, O., Ren, Z., Erkin, Z., & Lagendijk, R. L. (2018). Transaction propagation on permissionless blockchains: Incentive and routing mechanisms. In *2018 crypto valley conference on blockchain technology (cvcbt)* (pp. 20–30). IEEE. [↵](#)
88. Carlsten, M., Kalodner, H., Weinberg, S. M., & Narayanan, A. (2016). On the instability of bitcoin without the block reward. In *Proceedings of the 2016 acm sigsac conference on computer and communications security* (pp. 154–167). New York, NY, USA: ACM. <https://doi.org/10.1145/2976749.2978408> [↵](#)
89. Tsabary, I., & Eyal, I. (2018). The gap game. In *Proceedings of the 2018 acm sigsac conference on computer and communications security* (pp. 713–728). New York, NY, USA: ACM. <https://doi.org/10.1145/3243734.3243737> [↵](#)

90. Hardin, G. (1968). The tragedy of the commons. *Science*, 162(3859), 1243–1248. [↵](#)
91. Bonneau, J. (2016). Why buy when you can rent? - Bribery attacks on bitcoin-style consensus. In fcw16ed (Ed.), *Fcw16name* (Vol. fcw16vol, pp. 19–26). fcw16addr: fcwpub. [↵](#)
92. Liao, K., & Katz, J. (2017). Incentivizing blockchain forks via whale transactions. In *International conference on financial cryptography and data security* (pp. 264–279). Springer. [↵](#)
93. Velner, Y., Teutsch, J., & Luu, L. (2017). Smart contracts make bitcoin mining pools vulnerable. In *International conference on financial cryptography and data security* (pp. 298–316). Springer. [↵](#)
94. Teutsch, J., Jain, S., & Saxena, P. (2016). When cryptocurrencies mine their own business. In *International conference on financial cryptography and data security* (pp. 499–514). Springer. [↵](#)
95. McCorry, P., Hicks, A., & Meiklejohn, S. (2018). Smart contracts for bribing miners. *IACR Cryptology EPrint Archive*, 2018, 581. [↵](#)
96. Bonneau, J. (2018). Hostile blockchain takeovers (short paper). In *Bitcoin'18: Proceedings of the 5th workshop on bitcoin and blockchain research*. [↵](#)
97. Ritz, F., & Zugenmaier, A. (2018). The impact of uncle rewards on selfish mining in ethereum. In *2018 ieee european symposium on security and privacy workshops (euros pw)* (pp. 50–57). <https://doi.org/10.1109/EuroSPW.2018.00013> [↵](#)
98. Niu, J., & Feng, C. (2019). Selfish mining in ethereum. *ArXiv Preprint ArXiv:1901.04620*. [↵](#)
99. Buterin, V. (2016). *Uncle rate and transaction fee analysis*. Ethereum Foundation. Retrieved from <https://blog.ethereum.org/2016/10/31/uncle-rate-transaction-fee-analysis/> [↵](#)
100. Budish, E. (2018). *The economic limits of bitcoin and the blockchain*. National Bureau of Economic Research. [↵](#)
101. Sompolinsky, Y., & Zohar, A. (2015). Secure high-rate transaction processing in bitcoin. In *International conference on financial cryptography and data security*

- (pp. 507–527). Springer. [↵](#)
102. Sompolinsky, Y., Lewenberg, Y., & Zohar, A. (2016). *SPECTRE: A fast and scalable cryptocurrency protocol*. Cryptology ePrint Archive, Report 2016/1159. [↵](#)
 103. Sompolinsky, Y., & Zohar, A. (2018). *PHANTOM: A scalable blockdag protocol*. Cryptology ePrint Archive, Report 2018/104. [↵](#)
 104. *Snowflake to Avalanche: A Novel Metastable Consensus Protocol Family for Crypto*. (n.d.). [↵](#)
 105. Azouvi, S., McCorry, P., & Meiklejohn, S. (2018). Betting on blockchain consensus with fantomette. *CoRR, abs/1805.06786*. Retrieved from <http://arxiv.org/abs/1805.06786> [↵](#)
 106. Danezis, G., & Hrycyszyn, D. (2018). Blockmania: From block dags to consensus. *CoRR, abs/1809.01620*. Retrieved from <http://arxiv.org/abs/1809.01620> [↵](#)
 107. Kiayias, A., Russell, A., David, B., & Oliynykov, R. (2017). Ouroboros: A provably secure Proof of stake blockchain protocol. In *Annual international cryptography conference* (pp. 357–388). Springer. [↵](#)
 108. Pass, R., & Shi, E. (2017). Fruitchains: A fair blockchain. In *Proceedings of the acm symposium on principles of distributed computing* (pp. 315–324). ACM. [↵](#)
 109. Park, S., Pietrzak, K., Kwon, A., Alwen, J., Fuchsbauer, G., & Gaži, P. (2015). *SpaceMint: A cryptocurrency based on proofs of space*. Cryptology ePrint Archive, Report 2015/528. [↵](#)
 110. Bentov, I., Pass, R., & Shi, E. (2016). Snow White: Provably secure proofs of stake. *IACR Cryptology EPrint Archive, 2016*, 919. [↵](#)
 111. Buterin, V. (2017). *Incentives in casper the friendly finality gadget*. [↵](#)
 112. Fanti, G., Kogan, L., Oh, S., Ruan, K., Viswanath, P., & Wang, G. (2018). Compounding of wealth in proof-of-stake cryptocurrencies. *ArXiv Preprint ArXiv:1809.07468*. [↵](#)
 113. Kanjalkar, S., Kuo, J., Li, Y., & Miller, A. (2019). Short paper: I can’t believe it’s not stake! Resource exhaustion attacks on pos. In *International conference on financial cryptography and data security*. Springer. [↵](#)

114. Brown-Cohen, J., Narayanan, A., Psomas, C.-A., & Weinberg, S. M. (2018). Formal barriers to longest-chain proof-of-stake protocols. *ArXiv Preprint ArXiv:1809.06528*. [↵](#)
115. Lavi, R., Sattath, O., & Zohar, A. (2017). Redesigning bitcoin's fee market. *ArXiv Preprint ArXiv:1709.08881*. [↵](#)
116. Lavi, R. (2007). Computationally efficient approximation mechanisms. *Algorithmic Game Theory*, 301–329. [↵](#)
117. *Coin market cap*. (n.d.). Retrieved from <https://www.coinmarketcap.com/coins> [↵](#)
118. Gencer, A. E., Basu, S., Ittay Eyal and, R. van R., & Sirer, E. G. (2018). Decentralization in bitcoin and ethereum networks. In *Financial cryptography and data security*. Springer Berlin Heidelberg. [↵](#)
119. Arnosti, N., & Weinberg, S. M. (2018). Bitcoin: A natural oligopoly. *ArXiv Preprint ArXiv:1811.08572*. [↵](#)
120. Leonardos, N., Leonardos, S., & Piliouras, G. (2019). Oceanic games: Centralization risks and incentives in blockchain mining. *ArXiv Preprint ArXiv:1904.02368*. [↵](#)
121. Lewenberg, Y., Bachrach, Y., Sompolinsky, Y., Zohar, A., & Rosenschein, J. S. (2015). Bitcoin mining pools: A cooperative game theoretic analysis. In *Proceedings of the 2015 international conference on autonomous agents and multiagent systems* (pp. 919–927). Richland, SC: International Foundation for Autonomous Agents; Multiagent Systems. Retrieved from <http://dl.acm.org/citation.cfm?id=2772879.2773270> [↵](#)
122. Johnson, B., Laszka, A., Grossklags, J., Vasek, M., & Moore, T. (2014). Game-theoretic analysis of ddos attacks against bitcoin mining pools. In *International conference on financial cryptography and data security* (pp. 72–86). Springer. [↵](#)
123. Vasek, M., Thornton, M., & Moore, T. (2014). Empirical analysis of denial-of-service attacks in the bitcoin ecosystem. In *International conference on financial cryptography and data security* (pp. 57–71). Springer. [↵](#)
124. Laszka, A., Johnson, B., & Grossklags, J. (2015). When bitcoin mining pools run dry. In *International conference on financial cryptography and data security* (pp. 63–77). Springer. [↵](#)

125. Brünjes, L., Kiayias, A., Koutsoupias, E., & Stouka, A.-P. (2018). Reward sharing schemes for stake pools. *ArXiv Preprint ArXiv:1807.11218*. [↵](#)
126. Fiat, A., Koutsoupias, E., Ligett, K., Mansour, Y., & Olonetsky, S. (2013). Beyond myopic best response (in cournot competition). *Games and Economic Behavior*. [↵](#)
127. Luu, L., Velner, Y., Teutsch, J., & Saxena, P. (2017). SmartPool: Practical decentralized pooled mining. In *26th USENIX security symposium (USENIX security 17)* (pp. 1409-1426). Vancouver, BC: USENIX Association. Retrieved from <https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/luu> [↵](#)
128. Miller, A., Kosba, A., Katz, J., & Shi, E. (2015). Nonoutsourcable scratch-off puzzles to discourage bitcoin mining coalitions. In *Proceedings of the 22Nd acm sigsac conference on computer and communications security* (pp. 680-691). New York, NY, USA: ACM. <https://doi.org/10.1145/2810103.2813621> [↵](#)
129. Kwon, Y., Liu, J., Kim, M., Song, D., & Kim, Y. (2019). Impossibility of full decentralization in permissionless blockchains. *ArXiv Preprint ArXiv:1905.05158*. [↵](#)
130. Mccorrey, P., Möser, M., Shahandasti, S. F., & Hao, F. (2016). Towards bitcoin payment networks. In *Proceedings, part i, of the 21st australasian conference on information security and privacy - volume 9722* (pp. 57-76). New York, NY, USA: Springer-Verlag New York, Inc. [↵](#)
131. Miller, A., Bentov, I., Kumaresan, R., & McCorry, P. (2017). Sprites: Payment channels that go faster than lightning. *CoRR, abs/1702.05812*. Retrieved from <http://arxiv.org/abs/1702.05812> [↵](#)
132. Dziembowski, S., Eckey, L., Faust, S., & Malinowski, D. (2017). Perun: Virtual payment hubs over cryptocurrencies. In *Perun: Virtual payment hubs over cryptocurrencies* (p. 0). IEEE. [↵](#)
133. Dziembowski, S., Faust, S., & Hostáková, K. (2018). General state channel networks. In *Proceedings of the 25th acm sigsac conference on computer and communications security*. ACM. Retrieved from <https://eprint.iacr.org/2018/320> [↵](#)
134. Poon, J., & Dryja, T. (2016). *The bitcoin lightning network: Scalable off-chain instant payments*. Retrieved from <https://lightning.network/lightning-network->

paper.pdf [↵](#)

135. McCorry, P., Bakshi, S., Bentov, I., Miller, A., & Meiklejohn, S. (2018). *Pisa: Arbitration outsourcing for state channels*. Cryptology ePrint Archive, Report 2018/582. [↵](#)

136. Chen, X., Papadimitriou, C., & Roughgarden, T. (2019). An axiomatic approach to block rewards. In *Proceedings of the 1st acm conference on advances in financial technologies* (pp. 124–131). ACM. [↵](#)

137. Ford, B., & Böhme, R. (2019). Rationality is self-defeating in permissionless systems. *ArXiv Preprint ArXiv:1910.08820*. [↵](#)

138. Box, G. E. (1976). Science and statistics. *Journal of the American Statistical Association*, 71(356), 791–799. [↵](#)

139. Becker, G. S. (1962). Irrational behavior and economic theory. *Journal of Political Economy*, 70(1), 1–13. Retrieved from <http://www.jstor.org/stable/1827018> [↵](#)

140. Hern, A. (2014). Bitcoin currency could have been destroyed by 51% attack. *The Guardian*. Retrieved from <https://www.theguardian.com/technology/2014/jun/16/bitcoin-currency-destroyed-51-attack-ghash-io> [↵](#)

141. Halpern, J. Y., & Pass, R. (2013). Sequential equilibrium in computational games. In *IJCAI* (pp. 171–176). [↵](#)