



Rationals vs Byzantines in Consensus-based Blockchains

Yackolley Amoussou-Guenou, Bruno Biais, Maria Potop-Butucaru, Sara Tucci-Piergiovanni

► To cite this version:

Yackolley Amoussou-Guenou, Bruno Biais, Maria Potop-Butucaru, Sara Tucci-Piergiovanni. Rationals vs Byzantines in Consensus-based Blockchains. [Research Report] CEA List; LIP6, Sorbonne Université, CNRS, UMR 7606; HEC Paris. 2019. hal-02043331

HAL Id: hal-02043331

<https://hal.science/hal-02043331v1>

Submitted on 20 Feb 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Rationals vs Byzantines in Consensus-based Blockchains

Yackolley Amoussou-Guenou^{‡,*}, Bruno Biais[†],
Maria Potop-Butucaru^{*}, Sara Tucci-Piergiovanni[‡]

[‡]CEA, LIST, PC 174, Gif-sur-Yvette, 91191, France

^{*}Sorbonne Université, LIP6, CNRS, UMR 7606, Paris, France

[†]HEC Paris and Toulouse School of Economics, CNRS (TSM-Research), France

Abstract

In this paper we analyze from the game theory point of view Byzantine Fault Tolerant blockchains when processes exhibit rational or Byzantine behavior. Our work is the first to model the Byzantine-consensus based blockchains as a committee coordination game. Our *first* contribution is to offer a game-theoretical methodology to analyse equilibrium interactions between Byzantine and rational committee members in Byzantine Fault Tolerant blockchains. Byzantine processes seek to inflict maximum damage to the system, while rational processes best-respond to maximise their expected net gains. Our *second* contribution is to derive conditions under which consensus properties are satisfied or not in equilibrium. When the majority threshold is lower than the proportion of Byzantine processes, invalid blocks are accepted in equilibrium. When the majority threshold is large, equilibrium can involve coordination failures, in which no block is ever accepted. However, when the cost of accepting invalid blocks is large, there exists an equilibrium in which blocks are accepted iff they are valid.

1 Introduction

Since the publication of Nakamoto’s white paper [21] proposing the Proof-of-Work protocol, Bitcoin, thousands of blockchains have been created. At the operational level, a blockchain maintains an evolving list of ordered blocks. Each block consists of one or more transactions that have been verified by the system members. POW blockchains, however, consume excessive amounts of energy. This motivated tremendous efforts to propose alternatives protocols.

Byzantine-consensus based blockchains offer an alternative which has the advantage of being economical and offering strong consistency guarantees [4]. In Byzantine-consensus based blockchains such as HoneyBadger, HotStuff or Tendermint [1, 3, 11, 12, 17, 20, 25] a subset of deterministically selected processes, executes an instance of PBFT-consensus to decide on the next block to append. These protocols strive to satisfy the following properties: *Termination*: every non-Byzantine process decides on a value (a block); *Agreement*: if there is a non-Byzantine process that decides a value B , then all the non-byzantine processes decide B ; *Validity*[9]: a decided value by any non-Byzantine process is valid, it satisfies the predefined predicate.

While Byzantine consensus [18] is one of the best understood and formalized building blocks in distributed computing, blockchains systems revive this line of research in several respects: First, traditional Byzantine consensus has been analyzed only in systems where processes were either correct (verify their specification) or Byzantine (arbitrarily deviate from their specification). Blockchain systems bring on the scene a third type of player: rational players who take actions only if these actions increase their profits. Understanding the performance and limits of Byzantine-consensus

based blockchains with rational players is the goal of the current work. Our focus on rational players is in line with analyses of blockchain systems conducted by economists. Economists, however, have not considered Byzantine participants yet. Thus, our work endeavours to combine and unify computer science and economics approaches. Second, traditional Byzantine consensus analyses have not studied the choice and consequences of the way in which participants are rewarded. In this work we address the case of rewarding only the participants to the consensus.

Our contribution. Our contribution is twofold. First, we offer a methodology to analyse Byzantine consensus based blockchain protocols as a game between rational and Byzantine players. Two key aspects of the game, for rational players, are the cost of blocks verification and the cost of networking. Block verification is crucial since appending non verified blocks may have long term costs (e.g. double spending, collapse of the system etc.). Networking (participating to the agreement protocol by voting in favor of correct blocks) also has tremendous impact on system welfare: If participants don't vote, this can block the system or lead to agreement on invalid blocks. Second, we derive conditions (on the majority threshold necessary for block acceptance, ν , and the proportion of Byzantine processes, f) under which rational players reach an equilibrium where the consensus properties are guaranteed. Our findings are as follows. When $f \geq \nu$, invalid blocks are accepted, so that validity is not satisfied. When $f < \nu$, while there exists an equilibrium in which validity and termination are satisfied, there also exists an equilibrium in which blocks are never accepted, so that termination is not satisfied. This points to a tension between validity (which requires that the ν threshold be large enough) and termination (which can be threatened when the ν threshold is high.)

Related work. Blockchains can be roughly divided in consensus-less [21] and Byzantine consensus-based blockchains [17, 11, 1, 12, 3]. Byzantine Consensus-based blockchains have the advantage to guarantee strong consistency by running a Byzantine Fault Tolerant protocol [8]. In order to use a BFT protocol in an open setting, recent research has been devoted to either find secure mechanisms to select committees of fixed size over time (e.g. [16],[10]) and/or to propose incentives to promote participation [1]. Most of the proposals, however, assume participants as either honest or Byzantine, lacking to thoroughly explore the effect of rational participants. In this line of work, Solidus [1] is the first to consider rational processes by proposing an incentive-compatible BFT protocol for blockchains. Solidus introduces interesting incentive mechanisms, however, the paper lacks a game theoretic analysis of them.

While addressing a slightly different protocol, [19] is the closest work to ours. In this protocol multiple committees run in parallel to validate a non-intersecting set of transactions (a shard). A non-cooperative static game approach for the intra-committee protocol is taken leading to the result that rational agents can free-ride when rewards are equally shared. The main aspect of our analysis that is new and different from [19] is the following: we have a dynamic (not static) multi-round analysis, of a problem in which some participants are Byzantine and some blocks can be invalid (and costly for rational if accepted). In that context, there is a situation in which in equilibrium rational agents are pivotal, because if they do not check the block validity this will create the risk of having an invalid block accepted. It is because they are pivotal that they do not free ride. Moreover, we discuss equilibria in relation to formal consensus properties – Termination vs Validity –, which represents a novelty.

In the realm of consensus-less blockchains, as Bitcoin, many works used rational arguments to prove thresholds on the fraction of honest nodes needed to guarantee security properties [13, 24]. These works establish very pessimistic thresholds while in practice Bitcoin works even if the honest

majority assumption does not hold [5]. Following this observation, [5] proposes a rational analysis of Bitcoin based on the rational design protocol framework [15]. The proposed game, with respect to ours, is at an upper level of abstraction, proposing a two-player zero-sum game between the protocol designer and the adversary. Our game models instead the behavior of protocol participants, that can be rational, evolving in an environment with Byzantine processes. Moreover, our work targets consensus-based blockchain unlike [5].

With only rational players, [6] models Bitcoin as a coordination game. Similar to the work in [6], our analysis shows that the protocol in consensus-based blockchains is a coordination game. Additionally, we consider Byzantine players, and show that *Termination* can be violated when coordination failures occur. [23] uses a game theoretic approach to study consensus-less Proof-of-Stake Blockchains, and shows that the Nothing at Stake problem is mitigated because players with large stakes on the main chain prefer not to add blocks on forking branches, lest it should reduce the strength of the main chain, and thus the value of their stakes. The environment considered in [23] differs from ours, since the study in [23] does not consider consensus-based blockchains, nor Byzantine players.

2 Blockchain Consensus with Rational Players

2.1 System Model

We consider a system composed of a finite and ordered set Π , called *committee*, of synchronous sequential processes or players, namely $\Pi = \{p_1, \dots, p_n\}$ where process p_i is said to have index i . In the following, we refer to process p_i by its index, say process i . Hereafter, the words “player” and “process” are taken to have the same meaning.

Communication. We assume that each process evolves in rounds. A *round* consists of one or more phases, and each phase is divided into three sequential steps, in order: the send, the delivery and the compute step. We assume that the send step is atomically executed at the beginning of the phase and the compute step is atomically executed at the end of the phase. The phase has a fixed duration that allows collecting all the messages sent by the processes at the beginning of the phase during the delivery step. At the end of a phase a process exits from the current phase and starts the next one. The processes communicate by sending and receiving messages through a broadcast primitive. Messages are created with a digital signature, and we assume that digital signatures cannot be forged. When a process i delivers a message, it knows the process j that created the message. We assume that messages cannot be lost.

Processes Behavior. In this paper we consider a variant of the BAR model [2] where processes are either *rational* or *Byzantine*. *Rational processes* are self-interested and seek to maximize their expected utility. They will deviate from a prescribed (suggested) protocol if and only if doing so increases their expected utility. Their objective function must account for their costs (e.g., sending messages) and benefits (e.g., reward of a block) for participating in a system. In line with [2], the objective of Byzantine processes is prevent the protocol from achieving its goal, and to reduce the rational processes utility, no matter the cost. We denote by f the number of Byzantine processes in the network. We assume symmetric Byzantines, their behaviour is perceived identically by all non-Byzantine processes. That is, a message sent by a Byzantine process and received by a non-Byzantine process in a given phase is received by all non-Byzantine processes in the same phase.

2.2 Byzantine Consensus based Blockchain

Consensus-based blockchains should satisfy the following consensus properties:

- **Termination:** every non-Byzantine process decides on a value (a block);
- **Agreement:** if there is a non-Byzantine process that decides a value B , then all the non-byzantine processes decide B ;
- **Validity[9]:** a decided value by any non-Byzantine process is valid, it satisfies the predefined predicate.

Let us note that the above properties must hold also for systems prone to rational behavior.

To implement the above specification in Consensus-based blockchains, for each height $h > 0$ of the blockchain, a Consensus instance is run inside a committee selected for the given height. In this paper we analyze a very general protocol, inspired by [1, 3, 11, 12, 17, 20, 25], a variant of PBFT. In this protocol, a proposer proposes a value, i.e. a block, and the other members of the committee will check the validity of the value. If the value is valid, then they will vote for it and will announce their vote through a message to the other members. Votes are collected and if a given threshold is reached, then the value is decided, otherwise a new proposer will propose another block and the procedure restarts.

In this work, we study a protocol (for rational players) which in some equilibria implements the consensus. For the sake of clarity, we first present a prescribed protocol, and then the actions of the rational processes. If a rational player does not deviate from a given prescribed protocol, we can consider it as a correct process.

The prescribed protocol. The protocol proceeds in rounds. For sake of simplicity we consider the height k of the blockchain passed as parameter to the protocol. Algorithm 1 presents the pseudo-code of the protocol.

For each round t a committee member is designated as the proposer for the round in a round robin fashion. The `isProposer(t, k)` function returns true only if the process is the proposer for the current round (line 7). The function, by taking as parameter the current height, only returns true if the proposer is part of the current committee, deterministically selected on the basis on the information contained in the blockchain up to k (the actual selection mechanism is out of the scope of the paper). Each round is further divided in two phases: the PROPOSE and the VOTE phase.

During the PROPOSE phase, the proposer of the round uses the function `createValidValue(k)` to generate a block. Because a valid block must include the identifier of the k^{th} block in the blockchain, the height k is passed as parameter (line 8). Once the block is created, a message broadcasting the proposal is sent (line 9). At line 10 the proposal is received through a delivery function. Each process checks if the proposal is a valid value (line 13). If so, the process sets its vote to the value (line 14).

During the VOTE phase, any process that set its vote to the current valid proposal sends a message (of type vote) to the other members of the committee (line 18). During the delivery step, sent messages are collected by any process. During the compute step each process verifies if a quorum of ν votes for the current proposal has been reached. Let us note that ν , the majority threshold is a parameter here, because it is the object of our study to establish the quorum ν in presence of rational and Byzantine processes. If the quorum is reached, if the process voted for the value and did not already decided for the current height, then it decides for the current proposal (line 23) and the protocol ends. If the quorum is not reached, then a new round starts (line 26).

Algorithm 1 Prescribed Protocol for a given height k at any process i

```
1: Initialization:
2:    $vote := nil$ 
3:    $t := 1$ 
4:    $decidedValue := nil$ 
5: Phase PROPOSE( $t$ ):
6:   Send step:
7:   if  $i == isProposer(t, k)$  then
8:      $proposal \leftarrow createValidValue(k)$  /* The proposer of the round generates a block, i.e. the value to be proposed */
9:     broadcast (PROPOSE,  $k, t, proposal$ )
10:  Delivery step:
11:  delivery (PROPOSE,  $k, t, v$ ) from  $proposer(t)$  /* The process collects the proposal */
12:  Compute step:
13:  if  $isValid(v)$  then
14:     $vote \leftarrow v$  /* If the delivered proposal is valid, then the process sets a vote for it */
15: Phase VOTE( $t$ ):
16:  Send step:
17:  if  $vote \neq nil$  then
18:    broadcast (VOTE $i$ ,  $k, t, vote$ ) /* If the proposal is valid, the process sends the vote for it to all the validators */
19:  Delivery step:
20:  delivery (VOTE,  $k, t, v$ ) /* The process collects all the votes for the current height and round */
21:  Compute step:
22:  if  $|\langle VOTE, k, t, v \rangle| \geq \nu \wedge decidedValue = nil \wedge vote \neq nil \wedge vote = v$  then
23:     $decidedValue \leftarrow v$ ; exit /* The valid value is decided if the threshold is reached */
24:  else
25:     $vote \leftarrow nil$ 
26:     $t \leftarrow t + 1$ 
```

Let us note that the protocol in an environment assuming only correct (altruistic) and symmetric Byzantine processes trivially implements consensus if f , the number of Byzantine processes, is such that $f < \nu$. If $f \geq \nu$, on the other hand, the Termination property is not guaranteed. The scenario for that is that Byzantine validators might vote for a different value with respect to the one voted by correct processes or a nil value. In that case the correct process will not decide (line 22) and will move in the next round. The scenario can repeat forever.

In the following we detail the pseudo-code for a rational processes shown in Algorithm 2. The rational process will try to maximize its payoff by choosing to undertake or not the actions defined in its action space. We consider the choice of : (i) proposing or not a valid block, (ii) checking or not the validity of a block and (iii) sending or not the vote for a proposed block. The decision tree for the process i is shown in Figure 1.

Let us consider now rational processes. The rational process will try to maximize its payoff by choosing to undertake or not some actions, defined in Section 2.3. Intuitively, we consider the choice of : (i) proposing or not a valid block, (ii) checking or not the validity of a block and (iii) sending or not the vote for a proposed block. We consider that the actions of checking the validity of the block and the action of sending the message (of type vote) have a cost.

Protocol of the rational processes. Rational processes choices are explicitly represented in the pseudo-code (Algorithm 2) by dedicated variables, namely, $action^{propose}$, $action^{check}$, and $action^{send}$, defined at lines 5 – 7. Each action, initialized to nil , can take values from the set $\{0, 1\}$. Those values are set by calling the functions $\sigma_i^{propose}$, σ_i^{check} , and σ_i^{send} , respectively, returning the strategy for the process i .

Strategy $\sigma_i^{propose}$ determines if the proposer i chooses to produce a valid proposal or an invalid one (lines 12-16). In both cases the proposal is sent in broadcast (line 17).

Strategy σ_i^{check} determines if the receiving process chooses to check the validity of the proposal




Figure 1: Decision tree of process i

or not, which is a costly action. If the process chooses to check the validity (line 22), it will also update the knowledge it has about the validity of the proposal and it will pay a cost c_{check} . If otherwise, the process keeps not knowing if the proposal is valid or not ($validValue[t]$ remains set to \perp). Note that this value remains set to \perp even if the process is the proposer. This is because we assumed, without loss of generality, that checking validity has a cost and that the only way of checking validity is by executing the `isValid(v)` function.

Note that, as defined in Section 2.3, strategy σ_i^{send} depends on the knowledge the process has about the validity of the proposal. The strategy determines if the process chooses to send its vote for the proposal or not (line 24-30). If the process chooses to send a message for the proposal it will pay a cost c_{send} .

Let us note that the rational player that did not check the validity of the block could decide an invalid value if more than ν other processes have done the same and the proposed block is invalid.

We now define the game that represent the protocol.

2.3 Byzantine-Rational Game

Recall that out of the n players, $f \geq 1$ are Byzantine, while $n - f$ are rational. Each player i privately observes its own type, θ_i , which can be Byzantine ($\theta_i = \theta^b$) or rational ($\theta_i = \theta^r$).¹

Action space. As proposer, the player decides whether to propose a valid block or to propose an invalid block.

Then, at each round t , each player first decides whether to check the block's validity or not (at cost c_{check}), and second decides whether to send a message (at cost c_{send}) or not.

Information sets. At the beginning of each round $t > 1$, the information set of the player, h_i^t , includes the observation of the round number t , the player's own type θ_i , as well as the observation of what happened in previous rounds, namely (i) when the player decided to check validity, the knowledge of whether the block was valid or not, (ii) how many messages were sent, and (iii) whether a block was accepted or not. At round 1, h_i^1 only includes the player's private information about its own type, θ_i .

¹If player's type was observable (i.e., if Byzantine processes were detectable in advance) there would be a trivial solution to preclude them from harming the system: forbidding their participation.

Algorithm 2 Pseudo-code for a given height k modelling the rational process i 's behavior

```

1: Initialization:
2:    $vote := nil$ 
3:    $t := 1$  /* Current round number */
4:    $decidedValue := nil$ 
5:    $action^{propose} := nil$ 
6:    $action^{check} := nil$ 
7:    $action^{send} := nil$ 
8:    $validValue[] := \{\perp, \perp, \dots, \perp\}$  /*  $validValue[t] \in \{\perp, 0, 1\}$  */

9: Phase PROPOSE( $t$ ):
10:  Send step:
11:    if  $i == isProposer(k, t)$  then
12:       $action^{propose} \leftarrow \sigma_i^{propose}()$  /*  $\sigma_i^{propose} \in \{0, 1\}$  sets the action of proposing a valid block or an invalid one */
13:      if  $action^{propose} == 1$  then
14:         $proposal \leftarrow createValidValue(k)$ 
15:      else if  $action^{propose} == 0$  then
16:         $proposal \leftarrow createInvalidValue()$ 
17:      broadcast  $\langle PROPOSE, k, t, proposal \rangle$ 
18:  Delivery step:
19:    delivery  $\langle PROPOSE, k, t, v \rangle$  from proposer( $k, t$ )
20:  Compute step:
21:     $action^{check} \leftarrow \sigma_i^{check}()$  /*  $\sigma_i^{check} \in \{0, 1\}$  sets the action of checking or not the validity of the proposal */
22:    if  $action^{check} == 1$  then
23:       $validValue[t] \leftarrow isValid(v)$  /* The execution of  $isValid(v)$  has a cost  $c_{check}$  */
24:       $action^{send} \leftarrow \sigma_i^{send}(validValue)$  /*  $\sigma_i^{send} : \{\perp, 0, 1\} \rightarrow \{0, 1\}$  sets the action of sending the vote or not */
25:      if  $action^{send} == 1$  then
26:         $vote \leftarrow v$  /* The process decides to send the vote, the proposal might be invalid */

27: Phase VOTE( $t$ ):
28:  Send step:
29:    if  $vote \neq nil$  then
30:      broadcast  $\langle VOTE_i, k, t, vote \rangle$  /* The execution of the broadcast has a cost  $c_{send}$  */
31:  Delivery step:
32:    delivery  $\langle VOTE, k, t, v \rangle$  /* The process collects all the votes for the current height and round */
33:  Compute step:
34:    if  $|\langle VOTE, k, t, v \rangle| \geq \nu \wedge decidedValue = nil \wedge vote \neq nil \wedge vote = v$  then
35:       $decidedValue = v$ ; exit
36:    else
37:       $vote \leftarrow nil$ 
38:       $t \leftarrow t + 1$ 

```

Then, in each round $t > 1$, the player decides whether to check the validity of the current block. At this point, denoting by b_t the block proposed at round t , when the player does not decide to check validity $isValid(b_t)$ is the null information set, while if the player decides to check, $isValid(b_t)$ is equal to 1 if the block is valid and 0 otherwise. So, at this stage the player's information set becomes

$$H_i^t = h_i^t \cup isValid(b_t),$$

which is h_i^t augmented with the validity information player i has about b_t , the proposed block.

Strategies. At each round $t \geq 1$, the strategy of player i is a mapping from its information set into its actions. If the agent is selected to propose the block, its choice is given by $\sigma_i^{propose}(h_i^t)$. Then, at the point at which the agent can decide to check block's validity, its strategy is given by $\sigma_i^{check}(h_i^t)$. Finally, after making that decision, the player must decide whether to send a message or not, and that decision is given by $\sigma_i^{send}(H_i^t)$.

Reward and cost from adding blocks. In this paper we study the case in which, when a block is accepted, only the processes which sent a message are rewarded (and receive R). In addition, we

assume that when an invalid block is accepted, all rational players incur cost κ .

In this work we make the following assumption. The reward, R , to the players when a block is accepted is larger than the cost of checking validity, c_{check} , which in turn is larger than cost of sending, c_{send} , a message. But the reward obtained when a block is accepted is smaller than the cost of accepting an invalid block, κ . That is, $\kappa > R > c_{check} > c_{send}$.

Objective of rational players. Let T be the endogenous round at which the game stops. If a block is accepted at round $t \leq n$, then $T = t$. Otherwise, if no block is accepted, $T = n$. In the latter case, the *termination* property is not satisfied.

At the beginning of the first round, the expected gain of rational player i is:

$$U_i = E \left[\begin{aligned} & (R * \mathbb{1}_{(\sigma_i^{\text{send}}(H_i^T)=1)} * \mathbb{1}_{(\text{block accepted at } T)} - \kappa \mathbb{1}_{(\text{invalid block accepted})}) \\ & - \sum_{t=1}^T (c_{check} \mathbb{1}_{(\sigma_i^{\text{check}}(h_i^t)=1)} + c_{send} \mathbb{1}_{(\sigma_i^{\text{send}}(H_i^t)=1)}) \end{aligned} \middle| h_i^1 \right],$$

where $\mathbb{1}_{(.)}$ denotes the indicator function, taking the value 1 if its argument is true, and 0 if it is false.

Then, at the beginning of round $t > 1$, if $T \geq t$, the continuation payoff of the rational player with information set h_i^t is

$$W_{i,t}(h_i^t) = E \left[\begin{aligned} & (R * \mathbb{1}_{(\sigma_i^{\text{send}}(H_i^T)=1)} * \mathbb{1}_{(\text{block accepted at } T)} - \kappa \mathbb{1}_{(\text{invalid block accepted})}) \\ & - \sum_{s=t}^T (c_{check} \mathbb{1}_{(\sigma_i^{\text{check}}(h_i^s)=1)} + c_{send} \mathbb{1}_{(\sigma_i^{\text{send}}(H_i^s)=1)}) \end{aligned} \middle| h_i^t \right],$$

Objective of Byzantine players. In the current paper we assume the following: *Byzantine processes 1) as proposers, propose invalid blocks, and 2) when receiving a proposed block, check the blocks' validity and send a message if and only if the block is invalid.*

We conjecture, the above strategies will turn out to be the optimal strategies of the Byzantine players, minimizing $W_{i,t}$ in equilibrium.

Equilibrium concept. Since we consider a dynamic game, with asymmetric information, the relevant equilibrium concept is Perfect Bayesian Equilibrium [14], intuitively defined as follows:

A Perfect Bayesian equilibrium is such that all players 1) choose actions maximizing their objective function, 2) rationally anticipate the strategies of the others, and 3) draw rational inferences from what they observe, using their expectations about the strategies of the others and Bayes law, whenever it applies.

A Perfect Bayesian Equilibrium (PBE) is a Nash equilibrium [22], so players best-respond to one another. It imposes additional restrictions, to take into account the fact that the game is dynamic and that players can have private information, and therefore must draw rational inferences, from their observation of actions and outcomes. Rationality of inferences in PBE implies that (i) each player has rational expectations about the strategies of the others, and (ii) each player's beliefs are consistent with Bayes law, when computing probabilities conditional on events that have strictly positive probability on the equilibrium path. Perfection in PBE implies that, at each node starting a subgame the players' strategies form a Nash equilibrium of that subgame. In this context, to show that a strategy is optimal it is sufficient to show that it dominates any one-shot deviation [7].

Problem Definition In this work, we explore the behavior of rational players that could not validate the block – because checking validity has a cost – and conditions (the majority threshold ν and proportion of Byzantine processes) under which rational players reach an equilibrium where

consensus properties (defined in Section 2.2) are guaranteed. To do so, in Section 3, we study the equilibria that arise under different conditions.

3 Equilibria for Rational Players

3.1 Equilibrium when $f \geq \nu$

When the number of Byzantine players is larger than the majority rule, i.e., $f \geq \nu$, the validity property is not satisfied, since, when the first proposer is Byzantine, it proposes an invalid block, and that block is accepted, as all Byzantine players send messages in its favor. Against that backdrop, we characterize the strategies of the rational players and state the equilibrium outcome when $f \geq \nu$.

Proposition 1. *If $n - f \geq \nu + 1$ and $f \geq \nu$, there exists a Perfect Bayesian equilibrium in which the strategy of a rational player at any round is the following:*

- *As proposer, a rational player proposes a valid block.*
- *When receiving a proposed block, the rational players do not check the block validity but send a message.*

The first condition ($n - f \geq \nu + 1$) implies that, when all rational players but one send a message, they meet the majority threshold ν , so the block is accepted. The second condition ($f \geq \nu$) implies that, when all Byzantine processes send a message, the block is accepted. Under these conditions, each rational player understands it is not pivotal: If the block is invalid, Byzantine players will send messages, so that the block will be accepted irrespective of the rational player's own action. Moreover, if the block is valid, Byzantine players will not send messages, but all the other rational players will, so that the block will be accepted irrespective of the rational player's action.

Thus rational players understand that they are not pivotal, and that whatever they do, given the equilibrium behavior of the other rational agents and of the Byzantine processes; all blocks will be accepted. Consequently, they have no interest in checking the validity of the block. The only relevant comparison for them is between their expected gain when they send a message

$$R - c_{send} - \frac{f}{n}\kappa$$

and their expected gain when they do not send a message $-\frac{f}{n}\kappa$. Since, by assumption, $R > c_{send}$, rational players find it optimal to send a message. Finally note that, in the equilibrium of Proposition 1, a block is decided at round 1, so the *termination* property is satisfied, but, when the proposer is Byzantine, an invalid block is accepted, so the *validity* property is not satisfied.

Proof If a rational player is selected to be the proposer, he prefers to propose a valid block than to propose an invalid block. Indeed, if he proposes an invalid block, that block will be accepted (since the $f \geq \nu$ Byzantine players, on checking it and discovering it is invalid, will send a message). In that case the gain of the proposer is $R - c_{check} - c_{send} - \kappa$. If instead the rational player proposes a valid block, this block will be accepted and his gain will be $R - c_{check} - c_{send}$. Now, turn to the actions of rational players who are not proposers. The equilibrium gain of these players is

$$-c_{send} + R - \frac{f}{n}\kappa.$$

If instead of playing the equilibrium strategy, a rational player does not send a message, its expected gain is $-\frac{f}{n}\kappa$, which by assumption ($R > c_{send}$) is lower than the equilibrium expected gain.

Another deviation is to check the block's validity and send a message only if the block is valid, which brings expected gain equal to

$$-c_{check} + (1 - \frac{f}{n})(R - c_{send}) - \frac{f}{n}\kappa.$$

This is lower than the equilibrium expected gain if

$$-c_{send} + R - \frac{f}{n}\kappa > -c_{check} + (1 - \frac{f}{n})(R - c_{send}) - \frac{f}{n}\kappa,$$

which holds since it is equivalent to

$$0 > -c_{check} - \frac{f}{n}(R - c_{send}).$$

The other possible deviations are trivially dominated: Checking the block's validity and sending a message only when the block is invalid, yields expected gain

$$-c_{check} + \frac{f}{n}(R - c_{send} - \kappa),$$

which is lower than the equilibrium expected gain. Checking the block's validity and sending a message only when the block is valid yields expected gain

$$-c_{check} + \left(1 - \frac{f}{n}\right)(R - c_{send}) - \frac{f}{n}\kappa,$$

again lower than the equilibrium expected gain. Checks the validity of the block and always sending a message yields

$$R - c_{send} - \frac{f}{n}\kappa - c_{check},$$

which is again dominated, as is also checking and not sending, which yields $-c_{check} - \frac{f}{n}\kappa$.

□ *Proposition 1*

3.2 Equilibria when $f < \nu$

Proposition 2. *When $f < \nu$ and $n - f \geq \nu$, there exists a Nash equilibrium in which rational players never check blocks' validity nor send messages, so that no block is ever accepted.*

Condition $f < \nu$, in Proposition 2 implies that Byzantine players cannot reach the majority threshold on their own. This precludes accepting invalid blocks. So the *validity* property is satisfied. Unfortunately, the condition also implies there exists an equilibrium in which the *termination* property also fails to hold. The intuition is the following:

In Proposition 2, each rational player anticipates that no other player will send a message when the block is valid.² In this context, each rational player knows that, if it were to send a message in favor of a valid block, it would be the only one to do so. Because the majority threshold ν is strictly larger than 1, the block would not be accepted. Therefore sending a message is a dominated action

²Byzantine players send messages but only when the block is invalid.

for the rational player. Thus, the equilibrium in Proposition 2 reflects that rational players' actions are strategic complements and they must coordinate on sending messages in order to have valid blocks accepted. Proposition 2 shows that, in equilibrium, there can be a coordination failure, such that no block is ever accepted.³

Proof Consider a rational player who anticipates that other rational players will not send any message at any round. If it follow the equilibrium strategy and does not send an message, its gain is 0. This must be compared to the gain of the player if he deviates:

- If it sends a message without checking its expected gain is

$$-c_{send} + \Pr(\text{invalid})\mathbb{1}_{(f=\nu-1)}(R - \kappa).$$

- If it checks the block's validity and sends a message only when the block is valid, its expected gain is

$$-c_{check} - \Pr(\text{valid})c_{send}.$$

- If it checks the block's validity and sends a message only when the block is invalid, its expected gain is

$$-c_{check} + \Pr(\text{invalid})(\mathbb{1}_{(f=\nu-1)}(R - \kappa) - c_{send}).$$

- If it checks the validity of the block and always sends a message, its expected gain is

$$-c_{send} - c_{check} + \Pr(\text{invalid})\mathbb{1}_{(f=\nu-1)}(R - \kappa).$$

- If it checks and does not send a message, its gain is $-c_{check}$.

Clearly, the player is better off following the equilibrium strategy.

□ *Proposition 2*

Note that the conditions of Proposition 2 imply that $f < \frac{n}{2}$, i.e, there is a strict majority of rational players. Yet, the proposition shows that such majority is not enough to ensure both termination and validity.

While there exists an equilibrium in which termination does not obtain, this does not necessarily imply there is no equilibrium with termination and validity. To have termination and validity, it must be that, in equilibrium, sufficiently many rational players find it in their own interest to check the validity of the block and to send messages in support of valid blocks. The problem is that some players might be tempted to free-ride, and let the others bear the cost of checking. To avoid this situation, it must be that (at least some) rational players anticipate they are pivotal, i.e., if they fail to check block validity and send messages in support of valid blocks, this may derail the process at their own expense.

To make this point, we look for an equilibrium in which some rational players check the validity of the block and send a message if and only the block is valid, and this results in valid blocks being immediately accepted and invalid blocks being rejected. Before proving that such an equilibrium exists, we characterise the expected continuation payoff to which it would give rise.

³If $f = 0$, then, with $\nu = 1$, there exists a unique equilibrium, in which all processes check validity and send a message iff the block is valid. In that equilibrium validity and termination are satisfied. But this obtains only if there are no Byzantine processes. As soon as $f \geq 1$, if $\nu = 1$, Proposition 1 applies and validity is not satisfied.

Lemma 1. *Consider a candidate equilibrium in which some rational players check the validity of the block and send a message if and only the block is valid, while the other rational players send messages without checking validity, and this results in valid blocks being immediately accepted and invalid blocks being rejected. In such an equilibrium, if it exists, the expected continuation payoff, at round t , of the rational players who are to check block validity is*

$$\pi_{check}(t) = R - c_{send} - \phi(t)c_{check},$$

while the expected continuation payoff, at round t , of the rational players who are not to check block validity is

$$\pi_{send}(t) = R - \psi(t)c_{send},$$

where $\phi(f) = 1$, $\psi(f+1) = 1$ and both ϕ and ψ satisfy property P defined below.

Definition 1. *A function g satisfies property P , if $g(t) = 1 + \frac{f-t+1}{n-t+1}g(t+1), \forall t < f$.*

In the candidate equilibrium, participants will reach a point at which the block is valid and all rational players send a message so that the block accepted. This gives rise to a payoff $R - c_{send}$, the first part of $\pi_{check}(t)$. The second part of $\pi_{check}(t)$, $\phi(t)c_{check}$, is the expected cost of checking block validity, where $\phi(t)$ is the expected number of times the player expects to check validity before a block is accepted. Similarly, in $\pi_{send}(t)$, $\psi(t)c_{send}$, is the expected cost of sending messages, where $\psi(t)$ is the expected number of times the player expects to send messages before a block is accepted.

Proof We prove this Lemma in 2 parts:

1. Proof of the first part of the proposition, concerning the rational players who are expected to check validity:

At round $t = f$, players know that all $f - 1$ previous proposers were Byzantine and that there are now $n - f + 1$ potential proposers, out of which only one is Byzantine and $n - f > \nu$ are rational. The expected gains of the rational players who are supposed to check are

$$-c_{check} + \frac{n-f}{n-f+1}(R - c_{send}) + \frac{1}{n-f+1}(R - c_{send}),$$

where the first term is the cost of checking validity, the second term corresponds to the case in which the current proposer is rational and proposes a valid block that is immediately accepted, and the third term corresponds to the case in which the proposer is Byzantine, the block is rejected, and we move to the next round, at which a valid block is finally accepted (without needing any further validity check). This equilibrium payoff simplifies to

$$R - c_{send} - c_{check},$$

reflecting that eventually a valid block will be accepted, and that from round f on the player will need to check validity only once. This equilibrium payoff implies that

$$\phi(f) = 1.$$

Now turn to round $t < f$. If round $t \leq f$ is reached, the previous $t - 1$ proposers were Byzantine. There remains $n - (t - 1)$ potential proposers. Out of them a fraction

$$\frac{f - (t - 1)}{n - t + 1}$$

is Byzantine, while the complementary fraction

$$\frac{n-f}{n-t+1}$$

is rational. This fraction being the probability that the next proposer is rational.

To prove the property stated in the Proposition by backward induction, we now prove that if this property is satisfied at round $t+1$, that is if

$$\pi_{check}(t+1) = R - c_{send} - \phi(t+1)c_{check},$$

then it is satisfied at round t .

Suppose the rational player follows the equilibrium strategy of checking and sending iff the block is valid. Its expected gain from round t on is

$$-c_{check} + \frac{n-f}{n-t+1}(R - c_{send}) + \frac{f-(t-1)}{n-t+1}\pi(t+1),$$

where the first term is the cost of checking the block at round t , the second term is the probability that the block is valid and accepted multiplied by the payoff in that case, and the third term is the probability that the block is invalid and rejected multiplied by the payoff in that case. Substituting the value of $\pi_{check}(t+1)$, using that the property is verified at round $t+1$, the expected gain writes as

$$-c_{check} + \frac{n-f}{n-t+1}(R - c_{send}) + \frac{f-(t-1)}{n-t+1}(R - c_{send} - \phi(t+1)c_{check}).$$

That is

$$R - c_{send} - \left(1 + \frac{f-(t-1)}{n-t+1}\phi(t+1)\right) c_{check},$$

which, using the definition of $\phi(t)$, is $R - c_{send} - \phi(t)c_{check}$.

2. Proof of the second part of the proposition, concerning the rational players who are just expected to send messages:

Again, we prove that if the property is satisfied at round $t+1$, i.e., $\pi_{send}(t+1) = R - \psi(t+1)c_{send}$, then it is satisfied at round t . Suppose the rational player follows the equilibrium strategy of not checking blocks' validity and always sending a message. Its expected gain from round t on is

$$c_{send} + \frac{n-f}{n-t+1}R + \frac{f-t+1}{n-t+1}\pi_{send}(t+1),$$

where the first term is the cost of sending a message at round t , the second term is the probability that the block is valid and accepted multiplied by the payoff in that case, and the third term is the probability that the block is invalid and rejected multiplied by the payoff in that case. Substituting the value of $\pi_{send}(t+1)$, the expected gain writes as

$$-c_{send} + \frac{n-f}{n-t+1}R + \frac{f-t+1}{n-t+1}(R - \psi(t+1)c_{send}).$$

That is

$$R - \left(1 + \frac{f-t+1}{n-t+1}\psi(t+1)\right) c_{send},$$

which, using the definition of $\psi(t)$, is $R - \psi(t)c_{send}$.

□*Lemma 1*

Relying on Lemma 1, we now establish that our candidate equilibrium is indeed an equilibrium. To do so denote the highest index of all Byzantine players by i_B .

Proposition 3. *When $f < \nu$ and $n - f > \nu$, if the cost κ of accepting an invalid block is large enough, in the sense that*

$$\kappa > \alpha(t)c_{check} - \beta(t)c_{send}, \forall t < f,$$

where

$$\alpha(t) = \frac{(n - t + 1)\phi(t) - (f - t + 1)\Pr(i_B \geq n - \nu + f + 2 | T \geq t)\phi(t + 1)}{(f - t + 1)\Pr(i_B < n - \nu + f + 2 | T \geq t)}$$

and

$$\beta(t) = \frac{\Pr(i_B \geq n - \nu + f + 2 | T \geq t)}{\Pr(i_B < n - \nu + f + 2 | T \geq t)},$$

and if the reward is large enough relative to the costs in the sense that

$$R \geq \max \left[\frac{n}{n - f}c_{send}, c_{send} + \frac{n}{n - f}c_{check} \right],$$

there exists a Perfect Bayesian Nash equilibrium in which the strategy of rational players is the following:

- As proposer, a rational player proposes a valid block.
- At any round $t \leq f$, when receiving a proposed block, (i) the rational players with index $i \in \{t, \dots, n - \nu + f + 1\}$ check the block validity and send a message only if the block is valid, while (ii) the rational players with index $i \in \{n - \nu + f + 2, \dots, n\}$ do not check the validity of the block but send a message.
- If round $t = f + 1$ is reached, rational players send a message without checking if the block is valid. At this point the block is valid and accepted.

Hence, in equilibrium, termination occurs no later than at round $f + 1$.

On the equilibrium path, invalid blocks (proposed by Byzantine players) are rejected, while valid blocks (proposed by rational players) are accepted. This implies that, if round $t = f + 1$ is reached, the players know that during all the previous (f) rounds the proposers were Byzantine (to draw this inference, the rational players use their anticipation that all participants play equilibrium strategies; hence the Perfect Bayesian nature of the equilibrium). Consequently, at round $f + 1$, the proposer must be rational, and all players anticipate the proposed block is valid. So, no rational player needs to check the validity of the block but all send a message, which brings them expected gain equal to $R - c_{send}$. This is larger than their gain from deviating (e.g., by not sending a message or by checking the block.)

At previous rounds $t \leq f$, players know that all $t - 1$ previous proposers were Byzantine and that there remains $f - t + 1$ Byzantine players with index strictly larger than $t - 1$ (as above, this rational inference is a feature of the Perfect Bayesian equilibrium we characterize). Do the equilibrium strategies of the rational players preclude acceptance of an invalid block by Byzantine processes? To examine this point, consider the maximum possible number of messages that can be sent if the proposer is Byzantine. In equilibrium the $\nu - f - 1$ players with indexes strictly larger than $n - \nu + f + 1$ are to send a message without checking it. The worse case scenario

(maximizing the number of messages sent when the block is invalid) is that none of these players are Byzantine. In that case, in equilibrium, the number of messages sent when the block is invalid is $f + (\nu - f - 1) = \nu - 1$, so that we narrowly escape validation of the invalid block. In contrast, if one of the rational players deviated from equilibrium and sent a message without checking the block, in the worse case scenario, this would lead to accepting an invalid block. Thus, in that sense, the rational players with index strictly lower than $n - \nu + f + 1$ are pivotal. Hence they check block validity, because, under the condition stated in the proposition, the cost of accepting an invalid block is so large that rational players do not want to run that risk.

Proof For clarity, we decompose the proof in 5 steps.

1. The first step is to note that rational proposers strictly prefer to propose a valid block than an invalid one. This is because, when they follow their equilibrium strategy of proposing a valid block, it is accepted and the proposer gets $R - c_{check} - c_{send}$, while if they propose an invalid block, it is rejected, and we move to the next round, at which, in equilibrium, the player gets at most $R - c_{check} - c_{send}$ (and possibly less). Indeed, this player incurs the cost of checking validity at the next round, because the rational players who are not expected to check validity have indexes above $n - \nu + f + 1$, which are above $f + 1$, so that they do not get to propose blocks.
2. The next step concerns the actions of the rational players when round $t = f + 1$ is reached. At that round, all players know the proposer must be rational and the proposed block valid. In equilibrium no rational checks validity but all send a message. Any other action would be dominated.
3. The third step concerns the most relevant deviation, in which a rational player expected to check block validity fails to do so. If at round t a rational player supposed to check, deviates and sends a message without checking block validity, its expected continuation payoff is

$$\begin{aligned} & \left(1 - \frac{f - (t - 1)}{n - t + 1}\right) (R - c_{send}) + \frac{f - (t - 1)}{n - t + 1} \Pr(i_B < n - \nu + f + 1) (R - c_{send} - \kappa) \\ & + \frac{f - (t - 1)}{n - t + 1} \Pr(i_B \geq n - \nu + f + 1) (\pi(t + 1) - c_{send}). \end{aligned}$$

The first term is the payoff obtained by the deviating rational player if the current block is valid, and therefore immediately accepted. The second term is the payoff obtained by the deviating player when he was pivotal and triggered acceptance of an invalid block. To see this, consider the number of messages when the block is invalid, the rational player is deviating and the indexes of all the Byzantine players are strictly lower than $n - \nu + f + 2$: f messages are sent by the Byzantine processes, 1 message is sent by the deviating rational agent, $\nu - f - 1$ messages are sent by the rational players with index above than or equal to $n - \nu + f + 2$. The resulting total number of messages is ν and the block is accepted. The last term corresponds to the case in which the deviating rational player is not pivotal, and the invalid block is not accepted, so that we move to the next round.

Substituting the value of $\pi_{check}(t + 1) = R - c_{send} - \phi(t + 1)c_{check}$ from Lemma 1, the expected continuation value of the deviating player is

$$\begin{aligned} & \left(1 - \frac{f - (t - 1)}{n - t + 1}\right) (R - c_{send}) + \frac{f - (t - 1)}{n - t + 1} \Pr(i_B < n - \nu + f + 2 | T \geq t) (R - c_{send} - \kappa) \\ & + \frac{f - (t - 1)}{n - t + 1} \Pr(i_B \geq n - \nu + f + 2 | T \geq t) (R - c_{send} - \phi(t + 1)c_{check} - c_{send}). \end{aligned}$$

Or

$$(R - c_{send}) - \frac{f - (t - 1)}{n - t + 1} \Pr(i_B < n - \nu + f + 2 | T \geq t) \kappa \\ - \frac{f - (t - 1)}{n - t + 1} \Pr(i_B \geq n - \nu + f + 2 | T \geq t) (\phi(t + 1) c_{check} + c_{send}).$$

The equilibrium condition is that this deviation payoff must be lower than the equilibrium continuation payoff of the player

$$R - c_{send} - \phi(t) c_{check}.$$

That is

$$\frac{f - (t - 1)}{n - t + 1} \Pr(i_B < n - \nu + f + 2 | T \geq t) \kappa > \phi(t) c_{check} \\ - \frac{f - (t - 1)}{n - t + 1} \Pr(i_B \geq n - \nu + f + 2 | T \geq t) (\phi(t + 1) c_{check} + c_{send}).$$

Note that

$$\phi(t) \geq \frac{f - (t - 1)}{n - t + 1} \Pr(i_B \geq n - \nu + f + 2 | T \geq t) \phi(t + 1),$$

since by the definition of $\phi(t)$ this inequality is equivalent to

$$1 + \frac{f - (t - 1)}{n - t + 1} \phi(t + 1) \geq \frac{f - (t - 1)}{n - t + 1} \Pr(i_B \geq n - \nu + f + 2 | T \geq t) \phi(t + 1),$$

which indeed holds. Thus we can write the equilibrium condition as

$$\kappa > \alpha(t) c_{check} - \beta(t) c_{send}, \forall t < f,$$

as stated in the proposition.

4. Other possible deviations for rational player supposed to check block's validity are easier to rule out:

First, the player could do nothing (neither check nor send). Relative to the equilibrium payoff, this deviation economises the cost of checking (c_{check}). If the current proposer is Byzantine, the player then obtains the same payoff after a one shot deviation as on the equilibrium path ($\pi_{check}(t + 1)$). If the current proposer is rational, the block gets accepted, but the player does not earn any reward. So the deviation is dominated if

$$\frac{n - f}{n - t + 1} (R - c_{send}) \geq c_{check},$$

which holds under the condition, stated in the proposition, that $R \geq \max \left[\frac{n}{n-f} c_{send}, c_{send} + \frac{n}{n-f} c_{check} \right]$.

Second, the player could check the block validity, and then send a message irrespective of whether the block is valid or not. This would generate a lower payoff than the main deviation, shown above (in 3.) to be dominated.

Third, the player could check validity but then send no message. When the current proposer is Byzantine, this one-shot deviation yields the same payoff as the equilibrium strategy. When the current proposer is rational, this deviation yields a payoff of $-c_{check}$, which is lower than the equilibrium payoff $R - c_{send} - c_{check}$.

Fourth, the player could check the block's validity and send a message only if the block is invalid, which is trivially dominated.

5. Finally turn to deviations of rational players supposed to send messages without checking blocks' validity.

First, consider the possibility to abstain from sending a message. This economises the costs c_{send} , but, in case the block is valid and accepted, this implies the agent loses the reward R . So, the deviation is dominated if

$$\frac{n-f}{n-t+1}R \geq c_{send},$$

which holds under the condition, stated in the proposition, that $R \geq \max \left[\frac{n}{n-f}c_{send}, c_{send} + \frac{n}{n-f}c_{check} \right]$.

Second, consider the possibility of checking validity and sending a message only for valid blocks. This deviation would imply the agent would have to incur the cost of checking (c_{check}), but it would economise the cost of sending a message when the block is invalid. So the deviation is dominated if

$$c_{check} \geq \frac{f-t+1}{n-t+1}c_{send},$$

which holds because of our assumption that $c_{check} \geq c_{send}$.

Other deviations, such as checking validity but never sending messages, or checking validity and always sending messages, or checking validity and sending only if the block is invalid, are trivially dominated.

□ *Proposition 3*

4 Conclusion and Future Work

In this paper we model PBFT-consensus based blockchains as a coordination game between rational and Byzantine processes. We derive the conditions (on the majority threshold and the proportion of Byzantine processes) under which consensus properties are guaranteed in equilibrium or not. In future work, we will extend the analysis to more general Byzantine strategies and rational agents preferences, costs and rewards.

References

- [1] Ittai Abraham, Dahlia Malkhi, Kartik Nayak, Ling Ren, and Alexander Spiegelman. Solidus: An incentive-compatible cryptocurrency based on permissionless byzantine consensus. *CoRR*, abs/1612.02916v1, 2016.
- [2] Amitanand S. Aiyer, Lorenzo Alvisi, Allen Clement, Michael Dahlin, Jean-Philippe Martin, and Carl Porth. BAR fault tolerance for cooperative services. In *Proceedings of the 20th ACM Symposium on Operating Systems Principles 2005, SOSP 2005, Brighton, UK, October 23-26, 2005*, pages 45–58, 2005.
- [3] Yackolley Amoussou-Guenou, Antonella Del Pozzo, Maria Potop-Butucaru, and Sara Tucci-Piergiovanni. Correctness of tendermint-core blockchains. In *22nd International Conference on Principles of Distributed Systems, OPODIS 2018, December 17-19, 2018, Hong Kong, China*, pages 16:1–16:16, 2018.
- [4] Emmanuelle Anceaume, Antonella Del Pozzo, Romaric Ludinard, Maria Potop-Butucaru, and Sara Tucci-Piergiovanni. Blockchain Abstract Data Type. In *CoRR abs/1802.09877 and Poster at Proceedings of the 24th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, PPOPP 2019, Washington, DC, USA, February 16-20, 2019*, pages 439–440, 2019.
- [5] Christian Badertscher, Juan A. Garay, Ueli Maurer, Daniel Tschudi, and Vassilis Zikas. But why does it work? A rational protocol design treatment of bitcoin. In *Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part II*, pages 34–65, 2018.
- [6] Bruno Biais, Christophe Bisière, Matthieu Bouvard, and Catherine Casamatta. The blockchain folk theorem. *The Review of Financial Studies*, 2019.
- [7] David Blackwell. Discounted dynamic programming. *The Annals of Mathematical Statistics*, 36(1):226–235, 1965.
- [8] Miguel Castro and Barbara Liskov. Practical byzantine fault tolerance. In *Proceedings of the Third USENIX Symposium on Operating Systems Design and Implementation (OSDI), New Orleans, Louisiana, USA, February 22-25, 1999*, pages 173–186, 1999.
- [9] T. Crain, V. Gramoli, M. Larrea, and M. Raynal. (Leader/Randomization/Signature)-free Byzantine Consensus for Consortium Blockchains. <http://csrg.redbellyblockchain.io/doc/ConsensusRedBellyBlockchain.pdf>, 2017.
- [10] Bernardo David, Peter Gazi, Aggelos Kiayias, and Alexander Russell. Ouroboros praos: An adaptively-secure, semi-synchronous proof-of-stake blockchain. In *Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part II*, pages 66–98, 2018.
- [11] C. Decker, J. Seidel, and R. Wattenhofer. Bitcoin Meets Strong Consistency. In *Proceedings of the 17th International Conference on Distributed Computing and Networking Conference (ICDCN)*, 2016.

- [12] I. Eyal, A. E. Gencer, E. Gün Sirer, and R. van Renesse. Bitcoin-NG: A Scalable Blockchain Protocol. In *13th USENIX Symposium on Networked Systems Design and Implementation, (NSDI)*, 2016.
- [13] Ittay Eyal and Emin Gün Sirer. Majority is not enough: bitcoin mining is vulnerable. *Commun. ACM*, 61(7):95–102, 2018.
- [14] Drew Fudenberg and Jean Tirole. Perfect bayesian equilibrium and sequential equilibrium. *Journal of Economic Theory*, 53(2):236 – 260, 1991.
- [15] Juan A. Garay, Jonathan Katz, Ueli Maurer, Björn Tackmann, and Vassilis Zikas. Rational protocol design: Cryptography against incentive-driven adversaries. In *54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013, 26-29 October, 2013, Berkeley, CA, USA*, pages 648–657, 2013.
- [16] Yossi Gilad, Rotem Hemo, Silvio Micali, Georgios Vlachos, and Nikolai Zeldovich. Algorand: Scaling byzantine agreements for cryptocurrencies. In *Proceedings of the 26th Symposium on Operating Systems Principles, Shanghai, China, October 28-31, 2017*, pages 51–68, 2017.
- [17] E. Kokoris-Kogias, P. Jovanovic, N. Gailly, I. Khoffi, L. Gasser, and B. Ford. Enhancing Bitcoin Security and Performance with Strong Consistency via Collective Signing. In *Proceedings of the 25th USENIX Security Symposium*, 2016.
- [18] L. Lamport, R. Shostak, and M. Pease. The byzantine generals problem. *ACM Transactions on Programming Languages and Systems*, 4(3):382–401, July 1982.
- [19] Mohammad Hossein Manshaei, Murtuza Jadliwala, Anindya Maiti, and Mahdi Fooladgar. A game-theoretic analysis of shard-based permissionless blockchains. *IEEE Access*, 6:78100–78112, 2018.
- [20] Andrew Miller, Yu Xia, Kyle Croman, Elaine Shi, and Dawn Song. The honey badger of bft protocols. Cryptology ePrint Archive, Report 2016/199, 2016.
- [21] S. Nakamoto. Bitcoin: A Peer-to-Peer Electronic Cash System. <https://bitcoin.org/bitcoin.pdf>, 2008.
- [22] John Nash. Non-cooperative games. *Annals of Mathematics*, 54(2):286–295, 1951.
- [23] Fahad Saleh. Blockchain Without Waste: Proof-of-Stake. SSRN Scholarly Paper ID 3183935, Social Science Research Network, Rochester, NY, January 2019.
- [24] Ayelet Sapirshtein, Yonatan Sompolinsky, and Aviv Zohar. Optimal selfish mining strategies in bitcoin. In *Financial Cryptography and Data Security - 20th International Conference, FC 2016, Christ Church, Barbados, February 22-26, 2016, Revised Selected Papers*, pages 515–532, 2016.
- [25] Maofan Yin, Dahlia Malkhi, Mike Reiter, Guy Gueta, and Ittai Abraham. HotStuff: BFT Consensus in the Lens of Blockchain. *CoRR*, abs/1803.05069v2, 2018.