A Theory of BAR Games

Allen Clement, Jeff Napper, Harry Li, Jean-Philippe Martin[‡], Lorenzo Alvisi, Mike Dahlin
University of Texas at Austin, [‡]Microsoft Research-Cambridge

Abstract

Distributed systems that span multiple administrative domains require protocols that tolerate both Byzantine and selfish nodes. This paper offers a theory that can be used to analyze such protocols. The theory systematically extends traditional game theory solution concepts through an *ex ante* analysis that incorporates a rational player's awareness of the possible presence of Byzantine players in the player's utility function. We illustrate our approach by modeling synchronous Terminating Reliable Broadcast as a game. We show that Dolev and Strong's Byzantine TRB protocol with message authentication is not a Nash equilibrium and that rational deviations from it may lead to violation of the TRB safety properties. We present a new TRB protocol with the same asymptotic complexity of Dolev-Strong and prove it to be a Nash equilibrium. Finally, we prove that *(k-t)* robustness, a recently proposed solution concept for games with Byzantine and rational players, cannot yield an equilibrium in games, such as our TRB game, that model systems where any node may crash and communication is necessary and incurs cost.

1 Introduction

This paper introduces a new model for constructing cooperative services that span multiple administrative domains (MADs). Cooperative MAD services are attractive because their diffused control structure, where no central authority controls all participating nodes, may yield services that are potentially less costly and more democratic than their more centralized counterparts. Unfortunately, current models fail to capture critical aspects of MAD systems.

Traditionally, nodes deviate from their specification because they are *broken* (e.g., due to bugs, hardware failures, configuration errors, or even malicious attacks). MAD systems add a new dimension: without a central administrator ensuring that unbroken nodes faithfully follow their assigned protocol, nodes may also deviate from their specification because they are *selfish* and intent on maximizing their own utility.

Byzantine Fault Tolerance (BFT) [5, 16, 19] handles the first class of deviations well. However, the Byzantine model classifies *all* deviations as faults and requires some bound on the number of faults in the system; this bound is untenable when all nodes may benefit by deviating from the protocol. Conversely, models based on traditional game theory [26] do account for rational behavior, but they are brittle: although they handle selfish deviations, they may be vulnerable to arbitrary disruptions if even a single broken node deviates from the expected rational behavior.

Several recent efforts have attempted to provide a solid foundation to model and build MAD services. In previous work [3, 17, 20], we introduced the BAR model, named after the three classes of nodes (Byzantine, Altruistic, and Rational) that it explicitly considers. Byzantine nodes can deviate arbitrarily from their specification for any reason. Altruistic nodes follow their specification faithfully. Rational nodes behave selfishly and deviate from a given protocol if and only if doing so improves their own utility. We demonstrate protocols for state machine replication [3] and gossip-based multicast [17] that provably maintain their usual properties when most or all of the nodes act selfishly and the remainder act maliciously or malfunction in arbitrary ways. We use these protocols to build practical MAD services for cooperative backup and media streaming. These results suggest that BAR can produce tractable and accurate models of a large class of practical services under reasonable assumptions, but they do not establish a formal theory of BAR games. Moscibroda et al. [22] express formally the cost of adding Byzantine players to a purely selfish environment. They quantify this cost for a simple *inoculation game* using a notion of Byzantine Nash equilibrium in which selfish players maximize

their worst case outcome, leaving as open questions the analysis of more complex games such as those that require communication among nodes. players. Other recent efforts introduce new solution concepts that accommodate irrationally generous or malicious behaviors [1, 8], but their usefulness for distributed systems is limited since, as we prove in this paper, they cannot be used to model systems in which (a) any node may crash, (b) communication incurs cost, and (c) communication is necessary.

In this paper, we establish a formally sound foundation for modeling realistic MAD services. Our approach is to leverage existing *solution concepts*, such as Nash equilibria or Pareto optimality, by providing a systematic method to adapt these concepts to accommodate Byzantine behavior. We enrich utility functions to include the model of Byzantine behavior on which rational players base their decisions.

We illustrate our approach by revisiting the classic synchronous Terminating Reliable Broadcast (TRB) problem [12]. First, we show that Dolev-Strong's optimal synchronous TRB solution [7] for a Byzantine model with message authentication [15] fails when players are rational: we give an alternate strategy that is advantageous to each node if all other nodes follow Dolev-Strong but that violates safety if all nodes deviate. Then, we develop a novel TRB protocol that is incentive compatible in the presence of up to f ($f \leq n-1$) Byzantine players with the remaining players being either rational or altruistic.

Hence, this paper contains four main contributions:

- A proof that traditional BFT protocols can break in a BAR environment.
- A proof that no protocol is *(k,t)*-robust [1] in a system where any node may crash and communication is necessary and incurs cost.
- A theory of BAR games that includes a systematic approach for extending traditional game theory solution concepts to accommodate Byzantine behavior.
- A specification of TRB as a game and an incentive-compatible synchronous protocol, BaN TRB, that solves TRB ¡¡¡¡¡¡¡ .mine in that game. To our knowledge, this is the first paper to solve TRB in an environment where some nodes may be Byzantine and the remaining may be Rational. ======= in that game. To our knowledge, BaN TRB is the first paper to solve TRB in an environment where some nodes may be Byzantine and the remaining may be Rational. ¿¿¿¿¿¿¿ .r689

Although we focus on TRB, the approach outlined in this paper is general and intended to be used for analyzing large scale distributed MAD services like [3, 17].

2 Communication games

This paper considers communication games in which players perform a distributed computation via message passing. A *game* Γ is a 3-tuple $(\mathcal{N}, \mathcal{S}_\mathcal{N}, \mathcal{U})$. Let $\mathcal{N} = \{1, \ldots, n\}$ be the set of players in the game. The set $\mathcal{S}_i$ contains all possible strategies for player i while $\mathcal{S}_\mathcal{C} = \times_{i \in \mathcal{C}} \mathcal{S}_i$ is the set of all possible strategies for all players in $\mathcal{C} \subseteq \mathcal{N}$. A strategy σ_i denotes the strategy for player i. A *strategy profile* $\vec{\sigma}_\mathcal{C}$ associates a strategy to each player in $\mathcal{C} \subseteq \mathcal{N}$. We use $\vec{\sigma}$ to denote $\vec{\sigma}_\mathcal{N}$. A *utility function* u_i defines the utility that player i receives when the game is played with a specified strategy profile. The set $\mathcal{U}$ contains all such utility functions. The strategy profile determines the *outcome*, $(\text{RES}, \text{TRACE})$, of the game where RES is the result of the game and TRACE is the trace of performed actions. In a communication game, TRACE represents the complete sequence of messages sent from player i to player j for all pairs of players $i \neq j$. The utility for player i is a function of the benefits received from RES and the costs incurred by sending messages in TRACE.

Mechanism design [9, 21, 24] addresses the problem of specifying a game Γ and strategy profile $\vec{\sigma}$ to implement a desired functionality $\mathcal{F}$. A mechanism implements $\mathcal{F}$ if (a) the result RES of playing Γ with $\vec{\sigma}$ is $\mathcal{F}$ and (b) $\vec{\sigma}$ is *incentive-compatible*: playing σ_i maximizes u_i for all rational players i. A *solution concept* is a rule that uses $\mathcal{U}$ and $\mathcal{S}_\mathcal{N}$ to define a game equilibrium. Common solutions concepts from the literature include Nash Equilibrium,

backwards induction, Pareto optimality, and subgame perfect equilibria [26]. *Stronger* solution concepts impose stricter requirements for a strategy profile to be incentive-compatible. Traditional solution concepts assume that every player acts rationally—identifying solution concepts that are applicable when some players behave irrationally is an important problem [1, 8, 22, 29]. In the rest of this section, we present a theory that incorporates Byzantine behavior into existing equilibria and discuss previous approaches that define new equilibria in the presence of Byzantine players.

2.1 BAR Game Theory

To analyze Byzantine behavior in a game, we must account for how Byzantine actions impact rational players' utilities. Since these utilities depend on the number of, composition of, or strategies employed by Byzantine players, utility functions should account for these three factors.

We augment traditional utility functions by defining a *Byzantine aware utility function* characterized by these factors: *size:* the number of Byzantine players,[1] *play:* the distribution of Byzantine players among $\mathcal{N}$, and *strat:* the strategy profile played by the Byzantine players. We define a Byzantine aware utility function as follows:

Definition 1 (Byzantine Aware Utility Function)**.** *Let t be the maximum number of Byzantine players expected in the system and $T = \{\mathcal{T} \subseteq \mathcal{N} : |\mathcal{T}| \leq t\}$. A* Byzantine aware utility function *is the utility function:*

$$\bar{u}_i(\vec{\sigma}) = \underset{x\in[0...t]}{\text{size}} \circ \underset{\mathcal{T}\,:\,|\mathcal{T}|=x}{\text{play}} \circ \underset{\vec{\tau}_{\mathcal{T}}\in\mathcal{S}_{\mathcal{T}}}{\text{strat}} \circ u_i(\vec{\sigma}_{\mathcal{N}-\mathcal{T}}, \vec{\tau}_{\mathcal{T}})$$

consisting of size *(a function over the expected distribution of the number of Byzantine players) applied to* play *(a function over the expected distribution of the identity of the Byzantine players) applied to* strat *(a function over the expected distribution of the actions of Byzantine players) applied to u_i (a traditional utility function).*

[1] In this paper we assume the threshold Byzantine model. Extension to other models [19, 31] is future work.

Risk-Averse Players and Nash Equilibria. In this paper we focus on a Byzantine aware utility function for risk-averse players. *Risk-averse* players identify the worst-case utility without restricting the actions of Byzantine players [3, 17, 20, 22]. The Byzantine aware utility function for risk-averse player i is:

$$\bar{u}_i(\vec{\sigma}) = \min_{x\in[0...t]} \circ \min_{\mathcal{T}\,:\,|\mathcal{T}|=x} \circ \min_{\vec{\tau}_{\mathcal{T}}\in\mathcal{S}_{\mathcal{T}}} \circ u_i(\vec{\sigma}_{\mathcal{N}-\mathcal{T}}, \vec{\tau}_{\mathcal{T}})$$

In this paper we employ the Nash equilibrium solution concept [23] to evaluate strategy profiles:

$$\forall i \in \mathcal{N}, \forall \phi_i \in \mathcal{S}_i : \bar{u}_i(\vec{\sigma}) \geq \bar{u}_i(\vec{\sigma}_{\mathcal{N}-\{i\}}, \phi_i)$$

Intuitively, a strategy profile is a Nash equilibrium if no player i can increase its utility by unilaterally deviating from the protocol. We consider any protocol that is a Nash equilibrium to be incentive-compatible.

A General Theory. Although the protocol definition and analysis in the rest of this paper focus on a specific Byzantine aware utility function (risk-averse) and solution concept (Nash equilibrium), our formulation represents a general theory. Our theory allows for different Byzantine aware utility functions to be applied to different solution concepts. For example, in some environments it may be rational for players to maximize their *expected utility* given some model of the probabilities of different scenarios; in such an environment a Byzantine aware utility function would be expressed as $\bar{u}_i(\vec{\sigma}) = \sum_{x\in[0...t]} Pr(x) \circ \sum_{\mathcal{T}:|\mathcal{T}|=x} Pr(\mathcal{T}) \circ \sum_{\vec{\tau}_{\mathcal{T}}\in\mathcal{S}_{\mathcal{T}}} Pr(\vec{\tau}_{\mathcal{T}}) \circ u_i(\vec{\sigma}_{\mathcal{N}-\mathcal{T}}, \vec{\tau}_{\mathcal{T}})$.

Just as we can use different Byzantine aware utility functions, we can also use different solution concepts. In this paper, we focus on Nash equilibria, although we can extend our framework to encompass other concepts. As an example, we can utilize Bayesian Nash Equilibrium, in which players are allowed to adjust their strategies based on observations $\mathcal{K}$ of previous instances, by instantiating an appropriate Byzantine aware utility function that leverages these observations [20]:

$$\bar{u}_r(\vec{\sigma}, \mathcal{K}) = \min_{(x\in[0,t])|\mathcal{K}} \circ \min_{(\mathcal{T}\in\mathcal{N})|\mathcal{K}} \circ \min_{(\vec{\tau}_{\mathcal{T}}\in\mathcal{S}_{\mathcal{T}})|\mathcal{K}} \circ u_i(\vec{\sigma}_{\mathcal{N}-\mathcal{T}}, \vec{\tau}_{\mathcal{T}}, \mathcal{K})$$

The flexibility to combine the range of Byzantine aware utility functions with the range of existing solution concepts allows us to trade accuracy in our model of rational players for tractability in the selected solution concept. This contrasts with previous approaches [1, 8, 22] that provide narrow point solutions that do not explicitly generalize to other solution concepts. Unfortunately, as we show in the next section, two of these concepts though tractable, are inapplicable to an important class of distributed protocols.

2.2 k-FTNE and (k,t)-Robustness

Eliaz [8] introduced the k fault-tolerant Nash equilibrium (k-FTNE) to guarantee that a certain strategy is in a player's best interest even if at most k players are Byzantine. Abraham et al. [1] generalize k-FTNE with the (k,t)-robust equilibrium to incorporate collusion among a group of rational players in addition to the Byzantine behavior of other players. Specifically, for (a) a given strategy profile $\vec{\sigma}$, (b) a coalition of players $\mathcal{C}$ of size at most k[2] following a coalition strategy profile $\vec{\phi}_{\mathcal{C}}$, and (c) a set of Byzantine players of size at most t following a Byzantine strategy profile $\vec{\tau}_{\mathcal{T}}$, no rational player i in the coalition can obtain better utility than when the coalition follows the given strategy profile $\vec{\sigma}_{\mathcal{C}}$. We give the formal definition from [1]:

Definition 2 ((k,t)-robust equilibrium)**.** *A strategy profile $\vec{\sigma} \in \mathcal{S}_{\mathcal{N}}$ is a (k,t)-*robust equilibrium *if for all $\mathcal{C},\mathcal{T} \subseteq \mathcal{N}$, $\mathcal{C} \cap \mathcal{T} = \emptyset$, $|\mathcal{C}| \leq k$, and $|\mathcal{T}| \leq t$, $\forall \vec{\tau}_{\mathcal{T}} \in \mathcal{S}_{\mathcal{T}}$, $\forall \vec{\phi}_{\mathcal{C}} \in \mathcal{S}_{\mathcal{C}}$, $\forall i \in \mathcal{C}$ we have $u_i(\vec{\sigma}_{\mathcal{N}-\mathcal{T}}, \vec{\tau}_{\mathcal{T}}) \geq u_i(\vec{\sigma}_{\mathcal{N}-(\mathcal{C}\cup\mathcal{T})}, \vec{\phi}_{\mathcal{C}}, \vec{\tau}_{\mathcal{T}})$*

Although a (k,t)-robust equilibrium provides an attractive set of strong properties, we show that it is inapplicable to a broad class of games that encapsulate important properties of fault tolerant systems. We call such games *fault-tolerant communication games* and characterize them with three properties: (**FT1**) any single player can crash without preventing the desired outcome, (**FT2**) communication has cost, and (**FT3**) direct communication between at least two players is necessary. Fault-tolerant communication games encapsulate the concerns of a broad class of distributed systems primitives, like TRB, where reducing classical message complexity and tolerating crashes are important concerns [18, 10] as well as large-scale applications where free-riding to reduce communication costs is a significant concern [11, 6, 3, 17].

Definition 3 (Fault tolerant communication game)**.** *Γ_{FT} is a* fault tolerant communication game *iff* **(FT1)** *$\forall i \in \mathcal{N}$, $\exists \vec{\sigma} \in \mathcal{S}$ such that $\vec{\sigma}$ achieves functionality $\mathcal{F}$ and in which i neither sends nor receives a message by playing σ_i* **(FT2)** *sending a message incurs non-zero cost; and* **(FT3)** *$\forall \vec{\sigma} \in \mathcal{S}_{\mathcal{N}}$ that achieve $\mathcal{F}$ $\exists a, b \in \mathcal{N}$ such that a sends a message to b.*

Abraham et al. [1] show a secret sharing game to be (k,t)-robust. However, the game does not fulfill the requirements of a fault tolerant communication game. **FT2** is implicitly violated by an assumption that benefit is based solely on the result of the protocol, independent of the steps taken to reach that result. Either **FT1** or **FT3** is violated by reliance on a distinguished entity, called a central mediator, that presides over all communication; either the mediator is a player that cannot crash (violating **FT1**) or the mediator is not a player and no pair of players communicates directly (violating **FT3**). We consequently argue that (k,t)-robustness is too strong for any real system that is a fault tolerant communication game. For example, TRB as described in Section 3 is a fault tolerant communication game.

Theorem 1. *There is no (k,t)-robust strategy profile that achieves $\mathcal{F}$ for a fault tolerant communication game when $k > 0$ and $t > 0$.*

Proof. We proceed by contradiction. Assume there exists a strategy profile $\vec{\sigma}$ that is (k,t)-robust and achieves $\mathcal{F}$ for a fault tolerant game Γ. By **FT3** there exist two nodes i and j such that i sends a message to j. A (k,t)-robust strategy profile has maximal utility regardless of the composition of $\mathcal{C}$ and $\mathcal{T}$. Let $i \in \mathcal{C}$ and $j \in \mathcal{T}$ for $|\mathcal{C}| = 1$, $|\mathcal{T}| = 1$, and $\mathcal{C} \cap \mathcal{T} = \emptyset$ as noted by the definition of (k,t)-robust. Further, let j follow the crash strategy ψ_j posited

[2] Eliaz considers up to k faults, implicitly assuming $|\mathcal{C}| = 1$.

by property **FT1** and i follow a strategy ϕ_i obtained from σ_i assigned in $\vec{\sigma}_{\mathcal{N}}$ by removing all messages that i sends to j. Since j does not send or receive any messages, the result($\vec{\sigma}_{\mathcal{N}-\{i,j\}-\mathcal{T}}, \phi_i, \psi_j, \vec{\tau}_{\mathcal{T}}$) = result($\vec{\sigma}_{\mathcal{N}}$, $\vec{\tau}_{\mathcal{T}}$). It follows from **FT2** that $u_i(\vec{\sigma}_{\mathcal{N}-\mathcal{T}}, \vec{\tau}_{\mathcal{T}}) < u_i(\vec{\sigma}_{\mathcal{N}-(\{i\}\cup\mathcal{T})}, \phi_i, \vec{\tau}_{\mathcal{T}})$, which contradicts our initial assumption that $\vec{\sigma}_{\mathcal{N}}$ was (k,t)-robust. □

3 TRB

We illustrate the analysis and design of incentive-compatible communication games using Terminating Reliable Broadcast (TRB)—a fundamental primitive in distributed computing. In our context of BAR games, a TRB instance has four properties: (**TRB1**: Validity) If the leader is non-Byzantine and broadcasts value v, then each non-Byzantine processes eventually delivers v; (**TRB2**: Integrity) Each non-Byzantine process delivers at most one value, and if it delivers $v \neq$ sender faulty (**SF**) then the leader broadcast v; (**TRB3**: Agreement) If a non-Byzantine process delivers v, then all non-Byzantine processes eventually deliver v; (**TRB4**: Termination) Each non-Byzantine process eventually delivers a value.

The infinite-horizon [26] TRB game, $\Gamma_{\text{TRB}} = (\mathcal{N}, \mathcal{S}_{\mathcal{N}}, \mathcal{U})$, consists of an infinite sequence of TRB instances[3] in which player i is instance x's leader, denoted *leader*^{x}, if and only if $(x \bmod n) + 1 = i$ where $x > 0$. A strategy σ_i describes the protocol that player i follows to determine its actions during each instance of TRB in Γ_{TRB}. We assume at most f Byzantine players with unknown identities and that players communicate over reliable synchronous links. We also assume authenticated messages [15], meaning that (a) players sign messages and (b) signatures are unforgeable. A strategy profile $\vec{\sigma}$ achieves functionality $\mathcal{F}_{TRB}$ if and only if playing Γ_{TRB} with $\vec{\sigma}$ results in an infinite sequence of instances that all satisfy **TRB1–4**.

In the game Γ_{TRB}, we define a rational player i's benefits when $\vec{\sigma}$ is played as follows. In instances k in which i is not leader, $\textit{benefits}_i^k(\vec{\sigma}) = \varpi$ if **TRB2–4** hold. In each instance l where i is the leader $\textit{benefits}_i^l(\vec{\sigma}) = \beta + \varpi$ if **TRB1–4** hold and i broadcast v, $\textit{benefits}_i^l(\vec{\sigma}) = \varpi$ if **TRB2–4** hold and **TRB1** does not. In all instances h, if any of **TRB2–4** do not hold then $\textit{benefits}_i^h = 0$. Note that we permit $\varpi = 0$ (only the leader achieves a benefit for any given instance, but all nodes eventually may achieve benefit since leadership rotates) or $\beta = 0$ (all nodes achieve benefit each instance, with no special benefit for the leader). We require some combination of β and ϖ to exceed the expected communication costs in order to entice rational players to participate.

If $\vec{\sigma}$ is played, let $\text{sent}_{i\to j}^k(\vec{\sigma})$ be the sequence of messages sent by player i to j during instance k and $\text{sent}_i^k(\vec{\sigma})$ be the sequence of all messages i sends during instance k. Player i incurs $c_{snd}(m)$ cost from sending each message $m \in \text{sent}_i^k(\vec{\sigma})$, and we assume that c_{snd}(m) is proportional to *length*(m). We do not incorporate other costs like storage and computation, which are future work. We define the costs for player i in instance k to be

$$\textit{costs}_i^k(\vec{\sigma}) = \sum_{m \in \text{sent}_i^k(\vec{\sigma})} c_{snd}(m)$$

We define i's utility in the the k-th TRB instance to be $u_i^k(\vec{\sigma}) = \textit{benefits}_i^k(\vec{\sigma}) - \textit{costs}_i^k(\vec{\sigma})$. We evaluate player i's utility for a game if $\vec{\sigma}$ is played using the *expected average utility* over all TRB instances [26]:

$$\hat{u}_i(\vec{\sigma}) = \lim_{m\to\infty} \sum_{k=1}^{m} \frac{u_i^k(\vec{\sigma})}{m}$$

We define a Byzantine aware utility function for risk-averse player i that identifies the worst case utility for any possible population of Byzantine players and selection of Byzantine strategies as follows:

$$\bar{u}_i(\vec{\sigma}) = \min_{t\in[0\ldots f]} \circ \min_{\mathcal{T}\,:\,|\mathcal{T}|=t} \circ \min_{\vec{\tau}_{\mathcal{T}}\in\mathcal{S}_{\mathcal{T}}} \circ\, \hat{u}_i(\vec{\sigma}_{\mathcal{N}-\mathcal{T}}, \vec{\tau}_{\mathcal{T}})$$

4 BFT $\not\Longrightarrow$ Incentive-Compatible

We now show that traditional BFT implementations of TRB are vulnerable to the tragedy of the commons [13]: if non-Byzantine players are rational, then each player has an incentive to defect, but if

[3] We assume 128 bits are sufficient to count the instances in our infinite horizon game. In practice, cosmological events become important considerations before 2^{128} instances can be completed.

all rational players defect then $\mathcal{F}_{TRB}$ is not met.

For concreteness, we consider Dolev and Strong's synchronous TRB protocol (D-S TRB) [7]. In our context, running D-S TRB corresponds to assigning rational players a strategy profile $\vec{\delta}$. Each instance of D-S TRB proceeds through $f+1$ rounds. A valid message for p_j in round i has the form $m = \langle \text{VALUE}, v, k \rangle_{p_1,\ldots,p_i}$ where v is a value, k an instance number, the signatures $p_1, \ldots, p_i$ come from players distinct from each other and from p_j, and p_1 is *leader*^{k}. In round one, the leader broadcasts a signed message m containing a value v. In subsequent rounds, a non-leader player signs and forwards any valid message m containing a previously unobserved value v to all players that have not signed a message containing v with the optimization that a player forwards at most two unique values. In the last round, each player delivers v if v was the only value observed and delivers **SF** otherwise.

Consider an alternate *lazy* strategy λ that rational nodes may adopt. The lazy strategy is identical to δ_r except that in round $i \leq f$, rather than sending valid message m to *all* players as in δ_r, player r playing λ_r sends m to $f+1-s$ players who, to r's knowledge, have not signed a valid message containing v where s is the number of players r has already observed to sign a message containing v. In round $f+1$, λ_r is identical to δ_r. Note that, if $f=0$ or $f+1 \leq n \leq f+2$, λ_r reduces to δ_r.

The following Theorem 2 says that a single rational player r continues to receive the benefits that come from fulfilling $\mathcal{F}_{TRB}$ despite adopting λ_r. Theorem 3 states that since r incurs lower costs with λ_r than with δ_r, a rational player can achieve higher benefit by unilaterally deviating from $\vec{\delta}$. Theorem 4 shows that if all rational nodes try to improve their utility by employing λ, no player receives benefit as Byzantine players can force **TRB3** to fail in every instance. Proofs of Theorems 2– 4 appear in the Appendix.

Theorem 2 (Lazy Safety and Liveness (**TRB1-4**))**.** *For all $\mathcal{T} \subseteq \mathcal{N}$, $|\mathcal{T}| \leq f$, and $\forall \vec{\tau}_{\mathcal{T}} \in \mathcal{S}_{\mathcal{T}}$, if $\vec{\sigma} = (\vec{\delta}_{\mathcal{N}-\mathcal{T}-\{r\}}, \lambda_i, \vec{\tau}_{\mathcal{T}})$ is played for Γ_{TRB} then* **TRB1-4** *are fulfilled.*

Theorem 3 (Not a Nash Equilibrium)**.** *Consider $\bar{u}_r$ defined for a risk averse rational player r: $\bar{u}_r(\vec{\delta}_{\mathcal{N}-\mathcal{T},\{r\}}, \lambda_r) > \bar{u}_r(\vec{\delta})$ when $n > f+2$.*

Theorem 4 (Failed Agreement)**.** *If all non-Byzantine players follow strategy profile $\vec{\lambda}_{\mathcal{N}-\mathcal{T}}$, $|\mathcal{T}| \leq f$, $f > 1$ and $n > f+2$ then* **TRB3** *can fail in all instances of Γ_{TRB}.*

The breach of safety shown in Theorem 4 highlights a key contribution of this paper: we address communication games in which sending messages incurs cost. This influences our theory and impacts protocol design. For example, in games where sending messages is free (i.e., **FT2** does not hold), a wide variety of protocols, including D-S TRB, are incentive-compatible.

Theorem 5 (Zero Cost D-S TRB)**.** *D-S TRB is (k,t)-robust if $c_{snd}(m) = 0$ for all messages m.*

Proof-sketch (see Appnedix). Since D-S TRB fulfills $\mathcal{F}_{TRB}$, $\vec{\sigma}$ maximizes benefit. Since $c_{snd}(m) = 0$, $\vec{\sigma}$ has zero cost which cannot be further reduced. □

5 BaN TRB

In this section, we present Byzantine aware Nash TRB (BaN TRB)—an incentive-compatible extension of D-S TRB for rational, risk-averse players. An instance k of BaN TRB proceeds as D-S TRB through a series of $f+1$ rounds where the leader, *leader*^{k}, is player i if and only if $(k \bmod n)+1 = i$.

BaN TRB is based on three key ideas: predictable message sequences, expensive dummy messages, and punishments of violators. First, BaN TRB specifies predictable message sequences to enable a receiver to detect deviations such as omitted messages. Second, to make message sequences predictable, BaN TRB specifies that dummy messages sometimes be sent. Dummy message are structured to be expensive so that rational nodes prefer sending useful messages whenever possible. Third, if a receiver detects that a partner has sent a message sequence that departs from an expected pattern, the receiver *shuns* the sender, preventing the sender from benefitting from such deviations. This punishment

message	content	c_{snd}
VAL_i	$\langle \text{VALUE}, k, v\rangle_{leader^k, s_2, \ldots, s_{i-1}, r}$	γ
$VAL_\perp$	$\langle \text{VALUE}, k, \perp_{\{1,2\}}\rangle_r$	γ
PNC_x	$\langle \text{PENANCE}, k, x\rangle_r$	κ_x

Table 1: Costs and contents of specific messages sent by player r in instance k.

mechanism is commonly referred to as ∞-tit-for-tat or grim trigger [4].

Table 1 describes messages used in BaN TRB. *Value* messages, denoted VAL_i in round i, correspond to the *valid* messages sent in D-S TRB. Additionally, the special values $\perp_1$ and $\perp_2$, sent in a $VAL_\perp$ dummy value message, are used to ensure that a player can send two value messages in every instance, whether or not the leader of that instance is faulty. BaN TRB also specifies a dummy *penance* message to discourage unilateral deviations. A penance message $m = \langle \text{PENANCE}, k, t\rangle_q$ is *well-formed* if it contains an instance number k, a size indication t, is signed by the sender q, and $c_{snd}(m) \geq \kappa_t$. A *valid* penance message, denoted PNC_t, is received in round $f+1$ of an instance k and is sent in two cases: if a VAL_1 message is not sent during k, or if a penance message was sent during instance $(k-n)$—the last instance that had the same leader as instance k.

Figure 1 shows the BaN TRB protocol. Each player r maintains a set, denoted $shun_r$, across all instances of BaN TRB containing players that r has observed to deviate, and the protocol specifies that r sends no messages to members of $shun_r$. We use $\vec{\rho}$ to denote the strategy profile given by BaN TRB and ρ_r to denote player r's strategy in $\vec{\rho}$. In the first round, $leader^k$ selects a value v and forms an appropriate VAL_1 message. The leader r then sends the message to all non-shunned players. All players that receive appropriate VAL_1 messages from the leader extract the corresponding values and prepare to relay them with the restriction that at most two distinct values sent by the (obviously deviating) leader are extracted. If a player $p \neq leader^k$ does not receive a VAL_1 message from the leader or has sent a penance message for $leader^k$ in a previous round, then p constructs an appropriate PNC_x message (to be sent later) where $0 \leq x < n$ is the number of distinct leaders up to and including instance k that have forced p to send a penance message.

```
1   Protocol initialization for process p:
2     shun_p := ∅
3     foreach a ∈ N
4       penDuring[a] := ∅ ; recvdSeq[a] := ∅

6   Initialization for process p in instance k > 0:
7     extracted := ∅ ; relay := ∅ ; penance := ∅
8     leader := k mod |N|

10  Round 1, for p = leader, and value v:
11    send ⟨VALUE, k, v⟩_p to q ∈ N − shun_p − {p}
12    extracted := {v}

14  Round 1, for p ≠ leader:
15    when receive m = ⟨VALUE, k, v⟩_leader
16      if leader ∉ penDuring[p] then
17        if v ∉ extracted ∧ |extracted| < 2 then
18          relay ∪= {m}
19          extracted ∪= {v}
20      recvdSeq[leader] ∪= {m}
21    if extracted = ∅ then
22      penDuring[p] ∪= {leader} ; shun_p ∪= {leader}
23      penance := {⟨PENANCE, k, |penDuring[p]|⟩_p}

25  Round i, 2 ≤ i ≤ f for p:
26    foreach m = ⟨VALUE, k, v⟩_{leader,...,s_{i−1}} ∈ relay
27      send ⟨m⟩_p to q ∈ N − shun_p − {p}
28    relay := ∅
29    when receive m = ⟨VALUE, k, v⟩_{leader,...,s_i}
30      if v ∉ extracted ∧ |extracted| < 2 then
31        relay ∪= {m}
32        extracted ∪= {v}
33      recvdSeq[s_i] ∪= {m}

35  Round f + 1, for p:
36    if f > 0 then
37      foreach m = ⟨VALUE, k, v⟩_{leader,...,s_f} ∈ relay
38        send ⟨m⟩_p to q ∈ N − shun_p − {p}
39      relay := ∅
40      if |extracted| < 2 then
41        send ⟨VALUE, k, ⊥_1⟩_p to q ∈ N − shun_p − {p}
42      if |extracted| < 1 then
43        send ⟨VALUE, k, ⊥_2⟩_p to q ∈ N − shun_p − {p}
44      if penance ≠ ∅ then
45        send m ∈ penance to q ∈ N − shun_p − {p}
46      when receive m = ⟨VALUE, k, v⟩_{leader,...,s_{f+1}}
47        if v ∉ extracted ∪ {⊥_1, ⊥_2} then
48          extracted ∪= {v}
49        recvdSeq[s_{f+1}] ∪= {m}
50      when receive m = ⟨PENANCE, k, t⟩_q
51        penDuring[q] ∪= {leader}
52        if t = |penDuring[q]| then
53          recvdSeq[q] ∪= {m}
54      foreach q ∈ N − shun_p − {p}
55        if recvdSeq[q] ∉ M^k_{q→p} then
56          shun_p ∪= {q}
57    if |extracted| = 1 then
58      deliver v ∈ extracted
59    else
60      deliver SF
61      shun_p ∪= {leader}
```

Figure 1: BaN TRB for instance $k > 0$.

In rounds 2 through f, each player p after receiving any VAL_{i-1} messages, forwards VAL_i messages containing distinct values received to other players with the restriction that at most two values are forwarded over all the rounds. If p does not receive at least one VAL_2 message from player $q \neq leader^k$ in

round 2, then p expects a penance message from q in round $f+1$. Intuitively, a penance message is a more costly replacement for the missing value message.

In the final round $f+1$, player p first sends VAL_{f+1} messages to other players for any VAL_f messages received such that p sends only two distinct values over all rounds. If p has not then sent two values, p sends $VAL_\perp$ messages containing the placeholder values $\perp_1$ or $\perp_2$ as needed to ensure that two distinct values are sent. Specifying exactly two value messages in all instances provides a predictable sequence of messages between players. If p created a PNC_x message in round 1, then p sends that message to other players. At the end of round $f+1$, if p has extracted exactly one value $v \notin \{\perp_1, \perp_2\}$ then p delivers v; otherwise p delivers **SF**. Additionally, p shuns any player that sent an unacceptable message sequence to p.

In order to distinguish deviations from $\vec{\rho}$ to implement shunning, we define the *message sequence* sent from r to p when $\vec{\sigma}$ is played through instance k: $\text{seq}^k_{r\to p}(\vec{\sigma}) = \bigcup_{h\in[1,k]} \text{sent}^h_{r\to p}(\vec{\sigma})$. A message sequence is acceptable if it could have been sent by a player r following ρ_r. Formally,

Definition 4 (Acceptable Message Sequence)**.** *A message sequence from r to j through instance k is* acceptable *if and only if that sequence is in the set:*

$$\mathcal{M}^k_{r\to j} = \bigcup_{\substack{\forall \mathcal{C} \subseteq \mathcal{N} - \{r,j\},\\ \forall \vec{\sigma}_\mathcal{C} \in \mathcal{S}_\mathcal{C}}} \text{seq}^k_{r\to j}(\vec{\rho}_{\mathcal{N}-\mathcal{C}}, \vec{\sigma}_\mathcal{C})$$

For simplicity, $\mathcal{M}_{r\to j} \equiv \mathcal{M}^\infty_{r\to j}$.

An acceptable message sequence for an execution of BaN TRB consists of two valid value messages for each instance and if no VAL_1 message is sent in an instance k, a penance message for every instance $i \geq k$ where $leader^i = leader^k$. Further, for any two penance messages $\langle \text{PENANCE}, j, x\rangle_q$ and $\langle \text{PENANCE}, k, y\rangle_q$, if $j < k$ then $x \leq y$.

As discussed above, BaN TRB closely resembles D-S TRB. Consequently, the following Theorem holds that BaN TRB, like D-S TRB, maintains **TRB1-4**. The proof appears in the Appendix.

Theorem 6. *For all $\mathcal{T} \subseteq \mathcal{N}$, $|\mathcal{T}| \leq f$, and $\forall \vec{\tau}_\mathcal{T} \in \mathcal{S}_\mathcal{T}$, let $\vec{\sigma} = (\vec{\rho}_{\mathcal{N}-\mathcal{T}}, \vec{\tau}_\mathcal{T})$.* $\mathbf{TRB1-4}$ *hold in all instances of Γ_{TRB} when $\vec{\sigma}$ is played.*

6 BaN TRB Analysis

The Byzantine aware utility of a strategy profile for the game Γ_{TRB} is defined over the average utility of an instance of the TRB protocol. We determine the average utility by characterizing strategy profiles according to a distinguished player's view of the game when the profile is played. The use of shunning and penance messages naturally divides a player's view of other nodes into three classes, friends, ex-friends, and enemies, that we use to classify strategies that a rational player might explore. A *friend* of r maintains good relations with r by sending acceptable message sequences, while *ex-friends* do not. If r plays ρ_r, then an ex-friend of r is in $shun_r$. Not all ex-friends are created equal, though. When a leader does not send a VAL_1 message, it not only becomes an ex-friend of r, but forces r to send penance messages to friends, incurring extra costs. We thus call this subset of ex-friends, *enemies*, because they impose extra overhead by their actions.

Definition 5 (Friends)**.** *The* friends *of a player r at instance k when $\vec{\sigma}$ is played are* $\text{F}^k_r(\vec{\sigma}) = \{p \in \mathcal{N} - \{r\} : \text{seq}^k_{r\to p}(\vec{\sigma}) \in \mathcal{M}^k_{r\to p} \wedge \text{seq}^k_{p\to r}(\vec{\sigma}) \in \mathcal{M}^k_{p\to r}\}$.

Definition 6 (Ex-friends)**.** *The* ex-friends *of a player r at instance k when $\vec{\sigma}$ is played are* $\text{X}^k_r(\vec{\sigma}) = \mathcal{N} - \text{F}^k_r(\vec{\sigma}) - \{r\}$.

For simplicity, our analysis further differentiates between Byzantine, $\text{X}^{\text{BYZ}}_r(\vec{\sigma}) \subseteq \text{X}_r(\vec{\sigma}) \cap \mathcal{T}$, and non-Byzantine, $\text{X}^{\text{NON}}_r(\vec{\sigma}) \subseteq \text{X}_r(\vec{\sigma}) - \mathcal{T}$, ex-friends.

Definition 7 (Enemies)**.** *The enemies of player r at instance k when $\vec{\sigma}$ is played are* $\text{E}^k_r(\vec{\sigma}) = \text{E}^{k-1}_r(\vec{\sigma}) \cup \{p = \text{leader}^k : \langle VALUE, k, v\rangle_p \notin \text{sent}^k_{p\to r}(\vec{\sigma})\}$ *where trivially,* $\text{E}^0_r(\vec{\sigma}) = \emptyset$.

For simplicity, $\text{F}_r(\vec{\sigma}) \equiv \text{F}^\infty_r(\vec{\sigma})$, $\text{X}_r(\vec{\sigma}) \equiv \text{X}^\infty_r(\vec{\sigma})$, and $\text{E}_r(\vec{\sigma}) \equiv \text{E}^\infty_r(\vec{\sigma})$. The following Lemmas describe useful properties of friends: (a) if two players play $\vec{\rho}$, they remain mutual friends and (b) a player r that plays ρ_r is either a friend or enemy of another player. Proofs appear in the Appendix.

Lemma 1. *Suppose two players p and q play ρ_p and ρ_q, respectively. For all $\vec{v}_{\mathcal{N}-\{p,q\}} \in \mathcal{S}_{\mathcal{N}-\{p,q\}}$, let $\vec{\sigma} = (\vec{\rho}_{\{p,q\}}, \vec{v}_{\mathcal{N}-\{p,q\}})$. $p \in \mathrm{F}_q(\vec{\sigma})$.*

Lemma 2. *Suppose player p plays ρ_p. For all $\vec{v}_{\mathcal{N}-\{p\}} \in \mathcal{S}_{\mathcal{N}-\{p\}}$, let $\vec{\sigma} = (\rho_p, \vec{v}_{\mathcal{N}-\{p\}})$. If $q \in \mathcal{N}$ and $p \notin \mathrm{F}_q(\vec{\sigma})$, then $p \in \mathrm{E}_q(\vec{\sigma})$.*

We identify a steady state behavior in the infinite Γ_{TRB} with respect to a player's ex-friends and enemies to simplify calculating the average expected utility.

Definition 8 (Steady state)**.** *A game execution with strategy profile $\vec{\sigma}$ is in the* steady state *at instance k if and only if $\mathrm{E}_r^k(\vec{\sigma}) = \mathrm{E}_r(\vec{\sigma}) \wedge \mathrm{X}_r^k(\vec{\sigma}) = \mathrm{X}_r(\vec{\sigma})$.*

Every game eventually reaches the steady state because both $\mathrm{X}_r^k(\vec{\sigma})$ and $\mathrm{E}_r^k(\vec{\sigma})$ are non-decreasing sets (as k grows) that are bounded in size by $\mathcal{N}$. The steady state condition holds true for an infinite suffix of the TRB game and thus determines the average expected utility. We define $\mathit{costs}_r(\vec{\sigma})$ and $\mathit{benefits}_r(\vec{\sigma})$ to be the cost and benefit received by r over n consecutive instances in the steady state of $\vec{\sigma}$.

In the remaining analysis, we consider only the case where $n > f+1$, $f > 0$ and address the corner cases where $n = f+1$ or $f = 0$ in the Appendix.

6.1 Utility of Playing ρ

To prove that $\vec{\rho}$ is a Nash Equilibrium for risk-averse players, we first place a lower bound on the Byzantine aware utility that a player expects from playing the recommended strategy $\vec{\rho}$; in the next section, we then show that the lower bound of $\vec{\rho}$ corresponds to an upper bound on the utility a player expects from unilaterally deviating from $\vec{\rho}$.

The aware utility $\bar{u}_r(\vec{\rho})$ identifies a rational player r's worst-case utility when every non-Byzantine player follows $\vec{\rho}_{\mathcal{N}-\{p\}-\mathcal{T}}$ and the Byzantine players follow arbitrary strategies $\vec{\tau}_{\mathcal{T}} \in \mathcal{S}_{\mathcal{T}}$. In the steady state, we establish a lower bound on r's benefit and an upper bound on r's cost as a function of the friends and enemies of r to calculate the worst-case utility of following $\vec{\rho}$.

Benefits. By proving BaN TRB is a Byzantine fault-tolerant TRB protocol, we show that r receives full benefit when $(\vec{\rho}_{-\mathcal{T}}, \vec{\tau}_{\mathcal{T}})$ is played. The following Lemma is thus direct from Theorem 6:

Lemma 3. *For all $\vec{\tau}_{\mathcal{T}} \in \mathcal{S}_{\mathcal{T}}$, $\mathcal{T} \subseteq \mathcal{N}$, $|\mathcal{T}| \leq f$, let $\vec{\sigma} = (\vec{\rho}_{\mathcal{N}-\mathcal{T}}, \vec{\tau}_{\mathcal{T}})$. $\forall r \notin \mathcal{T}$: $\mathrm{benefits}_r(\vec{\sigma}) = \beta + n\varpi$.*

Costs. The costs of each TRB instance are determined by the number of *PNC* and *VAL* messages r sends to its friends (ex-friends are in shun_p). In the steady state, the number of friends and enemies is constant. We thus define $\mathit{costs}_r(\vec{\sigma})$ over the number of friends and enemies of r. To simplify the discussion, we further define the *cost of friendship*, $\mathrm{C}(x, y)$ to be the cost to a player r of following ρ_r where $\forall \vec{v}_{\mathcal{N}-\{r\}} \in \mathcal{S}_{\mathcal{N}-\{r\}}$, $x = |\mathrm{F}_r(\rho_r, \vec{v}_{\mathcal{N}-\{r\}})|$, $y = |\mathrm{E}_r(\rho_r, \vec{v}_{\mathcal{N}-\{r\}})|$.

Specifically, $\mathrm{C}(x, y) = x(y\kappa_y + 2n\gamma)$ where $n = |\mathcal{N}|$. In n instances of TRB, BaN TRB specifies that a player r send to all players not in shun_r (that is, to x friends): (a) two value messages (costing $2n\gamma$) and (b) a penance message of size κ_y where y is the number of enemies of r each time a member of enemies is leader (costing $y\kappa_y$).

We structure *PNC* messages so the cost κ_y increases with y, the number of enemies, to guarantee that a player cannot reduce costs by gaining enemies. Although r saves the cost of sending (a) $y-1$ penance messages to $n-y$ players (each costing κ_{y-1}) and (b) $2n$ value messages to the new enemy (each costing γ), the savings are eroded by distributing these savings over the y new *PNC* messages to $n-y-1$ players (each costing κ_y). We thus define κ_y as

$$\kappa_y = \begin{cases} \frac{(n-y)(y-1)\kappa_{y-1}+2n\gamma}{y(n-y-1)}, & y \in [1, n-2] \\ 0, & \text{otherwise} \end{cases}$$

From the definitions of $\mathrm{C}(x, y)$ and κ_y, we provide a Lemma to describe their properties: (a) it costs more to keep the same set of friends while making more enemies; (b) it costs more to have a player as an enemy than it does to keep him as a friend; and (c) costs are trivially minimized by having no friends:

Lemma 4. *Let $x \in [0, n-1], y \in [0, n-x-2]$.*
(a) $x > 0 \Rightarrow \mathrm{C}(x, y) \leq \mathrm{C}(x, y+1)$.
(b) $x > 0 \Rightarrow \mathrm{C}(x, y) \leq \mathrm{C}(x-1, y+1)$.
(c) $x = 0 \Rightarrow \mathrm{C}(x, y) = 0$.

The next Lemma bounds the maximum costs using the cost of friendship in the steady state when rational players follow $\vec{\rho}$.

Lemma 5. *Let $n > f+1, f > 0$. For all $\mathcal{T} \subseteq \mathcal{N}$, $|\mathcal{T}| \leq f$, and $\forall \vec{\tau}_{\mathcal{T}} \in \mathcal{S}_{\mathcal{T}}$, $\forall r \notin \mathcal{T}$, let $\vec{\varphi} = (\vec{\rho}_{\mathcal{N}-\mathcal{T}}, \vec{\tau}_{\mathcal{T}})$.* $\text{costs}_r(\vec{\varphi}) \leq \mathrm{C}(n-f-1, f)$

Proof. It follows from Lemma 1 that $|\mathrm{F}_r(\vec{\varphi})| \geq n-f-1$ and $|\mathrm{E}_r(\vec{\varphi})| \leq f$. It follows from Lemma 4 that $\mathrm{C}(x,y)$ is maximized when $x = f$. Since $\mathrm{C}(x,y)$ is defined for ρ, $costs_r(\vec{\varphi}) \leq \mathrm{C}(n-f-1, f)$. □

Utility. Using the bounds on steady state benefit and cost we provide a lower bound on Byzantine aware utility.

Lemma 6. *Let $n > f+1$ and $f > 0$. $\forall r \in \mathcal{N}$:* $\bar{u}_r(\vec{\rho}) \geq \frac{(\beta+n\varpi)-\mathrm{C}(n-f-1,f)}{n}$

Proof. The Byzantine aware utility under the risk-averse rational model depends upon the worst-case average expected utility. Let $\vec{\varphi} = (\vec{\rho}_{\mathcal{N}-\mathcal{T}}, \vec{\tau}_{\mathcal{T}})$, $\forall \mathcal{T} \subseteq \mathcal{N}$, $|\mathcal{T}| \leq f$. The average expected utility for any $r \in \mathcal{N} - \mathcal{T}$ is determined by the costs and benefits of the steady state, leading to $\hat{u}_r(\vec{\varphi}) = \frac{benefits_r(\vec{\varphi}) - costs_r(\vec{\varphi})}{n}$. Substituting according to Lemmas 3 and 5, we obtain $\hat{u}_r(\vec{\varphi}) \geq \frac{(\beta+n\varpi)-\mathrm{C}(n-f-1,f)}{n}$ for any Byzantine behavior, which gives the specified aware utility. □

6.2 Utility of Deviating

We now show that there exists a *spiteful strategy* for Byzantine players to follow that places an upper bound on a rational player r's average expected utility, irrespective of r's unilateral deviation. This upper bound matches the lower bound for $\bar{u}_r(\vec{\rho})$ and demonstrates that BaN TRB is a Nash equilibria for risk-averse players.

We define the spiteful strategy $\vec{\varsigma}^r_{\mathcal{T}}$ such that Byzantine players follow $\vec{\rho}_{\mathcal{T}}$, but collude against r by inserting r into $shun_q$ for all $q \in \mathcal{T}$. We first show that spiteful players are enemies of r:

Lemma 7. *For all $\mathcal{T} \subseteq \mathcal{N}$, $|\mathcal{T}| \leq f$, and $\forall r \notin \mathcal{T}$, let $\vec{o} = (\sigma_r, \vec{\rho}_{\mathcal{N}-\mathcal{T}-\{r\}}, \vec{\varsigma}^r_{\mathcal{T}})$. $\mathcal{T} \subseteq \mathrm{E}_r(\vec{o})$*

Proof. Any player p playing ς^r_p sends no messages to r and is thus in $\mathrm{E}_r(\vec{\sigma})$ by definition. □

Benefits. We demonstrate an upper bound on the benefit of any unilateral deviation by r.

Lemma 8. *Let $n \geq f+1$, $f \geq 0$. For all $\mathcal{T} \subseteq \mathcal{N}$, $|\mathcal{T}| \leq f$ and $\forall \sigma_r \in \mathcal{S}r$, $\forall r \notin \mathcal{T}$, let $\vec{o} = (\sigma_r, \vec{\rho}_{\mathcal{N}-\mathcal{T}-\{r\}}, \vec{\varsigma}^r_{\mathcal{T}})$.* $\text{benefits}_r(\vec{o}) \leq \beta + n\varpi$

Proof. Direct from benefits defined for Γ_{TRB}. □

A Lemma is required for the special case of deviations that result in r having no friends.

Lemma 9. *Let $n \geq f+1$, $f > 0$. For all $\mathcal{T} \subseteq \mathcal{N}$, $|\mathcal{T}| \leq f$, and $\forall \sigma_r \in \mathcal{S}r$, $\forall r \notin \mathcal{T}$, let $\vec{o} = (\sigma_r, \vec{\rho}_{\mathcal{N}-\mathcal{T}-\{r\}}, \vec{\varsigma}^r_{\mathcal{T}})$. If $|\mathrm{F}_r(\vec{o})| = 0$ and $|\mathrm{X}^{\mathrm{NON}}_r(\vec{o})| > 0$, then* $\text{benefits}_r(\vec{o}) \leq \varpi$

Proof sketch (Complete proof in Appendix).
Without friends, r cannot learn the values proposed by other players and must deliver **SF**, violating **TRB3**. When r is leader, it does not send a value message to other players, violating **TRB1**. Finally, since all other players will deliver **SF** when r is leader, as a special case r can obtain ϖ by also delivering **SF**. □

Costs. We next derive a lower bound on r's cost when f Byzantine players follow the spiteful strategy and r pursues any unilateral deviation. For deviations that maintain a non-zero number of friends, the following Lemma bounds the minimum cost of deviation:

Lemma 10. *Let $n \geq f+1$, $f > 0$. For all $\mathcal{T} \subseteq \mathcal{N}$, $|\mathcal{T}| = f$ and $\forall \sigma_r \in \mathcal{S}r$, $\forall r \notin \mathcal{T}$, let $\vec{o} = (\sigma_r, \vec{\rho}_{\mathcal{N}-\mathcal{T}-\{r\}}, \vec{\varsigma}^r_{\mathcal{T}})$. If $|\mathrm{E}_r(\vec{o})| < n-1$ then* $\text{costs}_r(\vec{o}) \geq \mathrm{C}(n-f-1, f)$.

Proof. Lemmas 7 and 2 imply that $\mathrm{X}_r(\vec{o}) = \mathrm{E}_r(\vec{o})$ so that a player is either a friend or enemy. Lemma 4 rule (b) then states that $\mathrm{C}(x,y)$ is minimized for $\min(x)$. Given the lower bound of $|\mathrm{E}_r(\vec{o})|$ determined by $|\mathcal{T}| = f$, minimal costs are attained for $\mathrm{C}(n-f-1, f)$. □

For deviations by player r described by $\vec{o}$ that maintain zero friends, that is, where $\mathrm{F}_r(\vec{o}) = \emptyset$, we note that r is not required to send any messages so that trivially, $costs_r(\vec{o}) \geq 0$.

Utility. Using the bounds on benefit and cost in the steady state, we prove an upper bound on r's utility.

Lemma 11. *Let* $n > f + 1$ *and* $f > 0$. $\forall r \notin \mathcal{T}$, $\forall \sigma_r \in \mathcal{S}_r$: $\bar{u}_r(\vec{\rho}_{\mathcal{N}-\{r\}}, \sigma_r) \leq \max\{\frac{(\beta+n\varpi)-\mathrm{C}(n-f-1,f)}{n}, \frac{\varpi}{n}\}$

Proof. To find the Byzantine aware utility under the risk-averse rational model, we find the worst-case average expected utility. For all $\mathcal{T} \subseteq \mathcal{N} - \{r\}$, $|\mathcal{T}| = f$, let $\vec{o} = (\vec{\rho}_{\mathcal{N}-\mathcal{T}-\{r\}}, \sigma_r, \vec{\varsigma}_{\mathcal{T}}^{\,r})$. The average expected utility for any $r \in \mathcal{N} - \mathcal{T}$ is determined by the costs and benefits of the steady state, leading to $\hat{u}_r(\vec{o}) = \frac{\textit{benefits}_r(\vec{o}) - \textit{costs}_r(\vec{o})}{n}$.

Consider first the case where $|\mathrm{E}_r(\vec{o})| < n-1$. It follows from Lemma 10 that $\textit{costs}_r(\vec{o}) \geq \mathrm{C}(n-f-1, f)$. Finally, $\bar{u}_r(\vec{o}) \leq \frac{(\beta+n\varpi)-\mathrm{C}(n-f-1,f)}{n}$ using the upper bound on benefits provided by Lemma 8.

Assume $|\mathrm{E}_r(\vec{o})| = n - 1$. It follows from Lemma 9 that $\textit{benefits}_r(\vec{o}) \leq \varpi$ and as argued above, that $\textit{costs}_r(\vec{o}) \geq 0$. Hence, $\bar{u}_r(\vec{o}) \leq \frac{\varpi}{n}$. □

6.3 BaN TRB Is a Nash Equilibrium

We prove that $\vec{\rho}$ is a Byzantine aware Nash Equilibrium using the bounds on Byzantine aware utility proved in the previous sections. In the presence of Byzantine behavior, we show that the minimum expected utility of executing BaN TRB is the maximum expected utility of any unilateral deviation strategy profile.

Before we prove the theorem, we first discuss the assumption that the protocol is worth playing when everyone cooperates. A sufficient condition for participating is that a player expects the benefits of running the protocol (successful agreements and proposals) to exceed the cost of doing so (messages) for the proposed protocol ρ. We state this assumption as $\beta + (n-1)\varpi \geq \mathrm{C}(n-f-1, f)$.

Theorem 7. *Let* $n > f + 1$ *and* $f \geq 0$. *Using the risk-averse rational model,* $\vec{\rho}$ *is a Nash equilibrium if* $\beta + (n-1)\varpi \geq \mathrm{C}(n-f-1, f)$.

Proof. It suffices to show $\forall i \in \mathcal{N}, \forall \sigma_i \in \mathcal{S}i$, $\bar{u}_i(\vec{\rho}) \geq \bar{u}_i(\vec{\rho}_{\mathcal{N}-\{i\}}, \sigma_i)$. It follows from Lemma 6 that $\bar{u}_i(\vec{\rho}) \geq \frac{\beta+n\varpi-\mathrm{C}(n-f-1,f)}{n}$ and from Lemma 11 that $\bar{u}_i(\vec{\rho}_{\mathcal{N}-\{i\}}, \sigma_i) \leq \max\{\frac{\beta+n\varpi-\mathrm{C}(n-f-1,f)}{n}, \frac{\varpi}{n}\}$. By our assumption that $\beta + (n-1)\varpi \geq \mathrm{C}(n-f-1, f)$, $\bar{u}_i(\vec{\rho}) \geq \bar{u}_i(\vec{\rho}_{\mathcal{N}-\{r\}}, \sigma_i)$, completing the proof. □

Price of Byzantine Anarchy and Malice. The *Price of Byzantine Anarchy* ($PoB(f)$) quantifies the cost imposed by Byzantine players compared to the optimal strategy, whereas the *Price of Malice* ($PoM(f)$) quantifies the cost imposed by f Byzantine players compared to no Byzantine players [22]. When $f = 0$, we can solve TRB in a single round where the leader broadcasts value v to all participants. The average utility of following this optimal protocol when there are no Byzantine players is $\frac{\beta+n\varpi-(n-1)\gamma}{n}$. Theorem 8 states the Price of Byzantine Anarchy and Price of Malice for BaN TRB, which follow directly from the utility of the optimal TRB protocol and the best and worst case utilities of following BaN TRB.

Theorem 8. *Consider* Γ_{TRB} *instantiated to tolerate* f *failures. If* $n > f + 1$, $f > 0$ *and all non-Byzantine players play BaN TRB then*
(a) $PoB(f) = \frac{\beta+n\varpi-\mathrm{C}(n-f-1,f)}{\beta+n\varpi-(n-1)\gamma}$
(b) $PoM(f) = \frac{\beta+n\varpi-\mathrm{C}(n-f-1,f)}{\beta+n\varpi-2n(n-1)\gamma}$.

An interesting question that we leave to future work is quantifying the *Price of Selfishness*–the additional cost required to achieve functionality $\mathcal{F}$ in a fault tolerant communication game when non-faulty players may be rational.

7 Related Work

Both game theory [23, 26] and Byzantine fault tolerance [5, 7, 16, 19, 31] have been extensively studied in isolation. Recent work applies theory to fault free distributed systems [2, 9, 14, 25, 27, 28, 30], but only limited efforts addressed faulty environments.

Recent work attempts to bridge the gap between these approaches. Aiyer et al. [3] and Li et al. [17] focus on practical considerations of building systems under the BAR model. Moscibroda et al. [22] express formally the cost of adding Byzantine players to a purely selfish environment for a simple *inoculation game* in which rational and Byzantine

players choose between two strategies—inoculate or not—and in which there is no communication between nodes. Our goals differ in that we focus on showing that a specific strategy profile (a) meets a specified functionality $\mathcal{F}_{TRB}$ and (b) is also a Nash Equilibrium, while they focus on analyzing the additional costs Byzantine players are able to impose on the system.

Eliaz [8] defines k-FTNE to provide an appealing equilibrium concept in the presence of at most k faulty players. Eliaz analyzes an important problem in economics, the constrained Walrasian function.

Abraham et al. [1] generalize k-FTNE to (k, t)-robustness, which accomodates collusion in addition to faulty behavior, and analyze a protocol for resilient secret sharing with a centralized, failure-free, mediator with free communication.

8 Conclusion

In this paper we present a theory of BAR games that provides for the analysis of games when some subset of the players are irrational, and we analyze a novel protocol for Terminating Reliable Broadcast in the context of this theory.

References

[1] I. Abraham, D. Dolev, R. Gonen, and J. Halpern. Distributed computing meets game theory: robust mechanisms for rational secret sharing and multiparty computation. In *Proc. 25th PODC*, July 2006.

[2] M. Afergan and R. Sami. Repeated-game modeling of multicast overlays. In *IEEE INFOCOM 2006*, Apr. 2006.

[3] A. S. Aiyer, L. Alvisi, A. Clement, M. Dahlin, J.-P. Martin, and C. Porth. BAR fault tolerance for cooperative services. In *Proc. 20th SOSP*, Oct. 2005.

[4] R. Axelrod. *The Evolution of Cooperation*. Basic Books, New York, 1984.

[5] M. Castro and B. Liskov. Practical Byzantine fault tolerance and proactive recovery. *ACM TOCS*, Nov. 2002.

[6] B. Cohen. The BitTorrent home page. http://bittorrent.com.

[7] D. Dolev and H. R. Strong. Authenticated algorithms for Byzantine agreement. *Siam Journal Computing*, 12(4), Nov. 1983.

[8] K. Eliaz. Fault tolerant implementation. *Review of Economic Studies*, 69, Aug 2002.

[9] J. Feigenbaum and S. Shenker. Distributed algorithmic mechanism design: Recent results and future directions. In *Proc. 6th DIALM*. ACM Press, New York, 2002.

[10] P. Fraigniaud, C. Gavoille, and B. Mans. Interval routing schemes allow broadcasting with linear message-complexity. New York, 2000. ACM Press.

[11] Gnutella. http://www.gnutella.com/.

[12] V. Hadzilacos and S. Toueg. Fault-tolerant broadcasts and related problems. *Dist. Systems (2nd Ed.)*, 1993.

[13] G. Hardin. The tragedy of the commons. *Science*, 162, 1968.

[14] I. Keidar, R. Melamed, and A. Orda. Equicast: Scalable multicast with selfish users. In *Proc. 25th PODC*, 2006.

[15] L. Lamport. Time, clocks, and the ordering of events in a distributed system. *Comm. of the ACM*, 21(7), July 1978.

[16] L. Lamport, R. Shostak, and M. Pease. The Byzantine generals problem. *ACM Trans. Program. Lang. Syst.*, 4(3), 1982.

[17] H. C. Li, A. Clement, E. Wong, J. Napper, I. Roy, L. Alvisi, and M. Dahlin. BAR Gossip. In *Proc. 7th OSDI*, Nov. 2006.

[18] N. A. Lynch. *Distributed Algorithms*. Morgan Kaufmann Publishers Inc., San Francisco, CA, 1996.

[19] D. Malkhi and M. Reiter. Byzantine quorum systems. *Distributed Computing*, 11(4), 1998.

[20] J.-P. Martin. *Byzantine Fault-Tolerance and Beyond*. PhD thesis, UT Austin, Dec. 2006. TR-06-66.

[21] D. Monderer and M. Tennenholtz. Distributed games: from mechanisms to protocols. In *In Proc. 16th AAAI/11th IAAI*, Menlo Park, CA, 1999. AAAI.

[22] T. Moscibroda, S. Schmid, and R. Wattenhofer. When selfish meets evil: Byzantine players in a virus inoculation game. In *Proc. 25th PODC*, July 2006.

[23] J. Nash. Non-cooperative games. *The Annals of Mathematics*, 54, Sept 1951.

[24] N. Nisan and A. Ronen. Algorithmic mechanism design. In *In Proc. 31st STOC*, New York, NY, USA, 1999. ACM Press.

[25] N. Nisan and A. Ronen. Algorithmic mechanism design. *Games and Economic Behavior*, 35, April 2001.

[26] M. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, 1994.

[27] C. Papadimitriou. Algorithms, games, and the internet. In *Proc. 33rd STOC*. ACM Press, 2001.

[28] T. Roughgarden and E. Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2), 2002.

[29] J. Shneidman and D. C. Parkes. Specification faithfulness in networks with rational nodes. In *Proc. 23rd PODC*. ACM Press, 2004.

[30] E. Tardos. Network games. In *36th STOC*. ACM Press, 2004.

[31] P. Thambidurai and Y.-K. Park. Interactive consistency with multiple failure modes. In *Proc. 7th SRDS*, 1988.

```
1   Initialization for process p in instance k > 0:
2     leader := k mod |N|
3     extracted := ∅
4     relay := ∅

7   Round 1, for p = leader, and value v:
8     extracted := {v}
9     send ⟨VALUE, k, v⟩_p to q ∈ N − {p}

11  Round 1, for p ≠ leader:
12    when receive ⟨VALUE, k, v⟩_leader
13      if v ∉ extracted ∧ |extracted| < 2 then
14        relay ∪= {⟨VALUE, k, v⟩_leader}
15      extracted ∪= {v}

18  Round i, 2 ≤ i ≤ f for p:
19    foreach ⟨VALUE, k, v⟩_{leader,...,s_{i−1}} ∈ relay
20      send ⟨VALUE, k, v⟩_{leader,...,s_{i−1},p} to
              q ∈ N − sigs_v − {p}
21    relay := ∅
22    when receive ⟨VALUE, k, v⟩_{leader,...,s_i}
23      if v ∉ extracted ∧ |extracted| < 2 then
24        relay ∪= ⟨VALUE, k, v⟩_{leader,...,s_i}
25        sigs_v := {leader, ..., s_i}
26        extracted ∪= {v}
27      else if v ∈ extracted then
28        sigs_v ∪= {leader, ..., s_i}

30  Round f + 1 for p:
31    foreach ⟨VALUE, k, v⟩_{leader,...,s_f} ∈ relay
32      send ⟨VALUE, k, v⟩_{leader,...,s_f,p} to q ∈ N − sigs_v − {p}
33    relay := ∅
34    when receive ⟨VALUE, k, v⟩_{leader,...,s_{f+1}}
35      if v ∉ extracted ∧ |extracted| < 2 then
36        extracted ∪= {v}
37    if |extracted| = 1 then
38      deliver v ∈ extracted
39    else
40      deliver SF
```

Figure 2: Dolev-Strong protocol for instance $k > 0$.

```
1   Initialization for process p in instance k > 0:
2     leader := k mod |N|
3     extracted := ∅
4     relay := ∅

7   Round 1, for p = leader, and value v:
8     extracted := {v}
9     R ⊆ N − {p} : |R| = f + 1
10    send ⟨VALUE, k, v⟩_p to q ∈ R

12  Round 1, for p ≠ leader:
13    when receive ⟨VALUE, k, v⟩_leader
14      if v ∉ extracted ∧ |extracted| < 2 then
15        relay ∪= {⟨VALUE, k, v⟩_leader}
16      extracted ∪= {v}

19  Round i, 2 ≤ i ≤ f for p:
20    foreach ⟨VALUE, k, v⟩_{leader,...,s_{i−1}} ∈ relay
21      R ⊆ N − sigs_v − {p} : |R| =min(n − 1, f + 1) − |sigs_v|
22      send ⟨VALUE, k, v⟩_{leader,...,s_{i−1},p} to q ∈ R
23    relay := ∅
24    when receive ⟨VALUE, k, v⟩_{leader,...,s_i}
25      if v ∉ extracted ∧ |extracted| < 2 then
26        relay ∪= ⟨VALUE, k, v⟩_{leader,...,s_i}
27        sigs_v := {leader, ..., s_i}
28        extracted ∪= {v}
29      else if v ∈ extracted then
30        sigs_v ∪= {leader, ..., s_i}

32  Round f + 1 for p:
33    foreach ⟨VALUE, k, v⟩_{leader,...,s_f} ∈ relay
34      send ⟨VALUE, k, v⟩_{leader,...,s_f,p} to
              q ∈ N − sigs_v − {p}
35    relay := ∅
36    when receive ⟨VALUE, k, v⟩_{leader,...,s_{f+1}}
37      if v ∉ extracted ∧ |extracted| < 2 then
38        extracted ∪= {v}
39    if |extracted| = 1 then
40      deliver v ∈ extracted
41    else
42      deliver SF
```

Figure 3: Lazy variation of Dolev-Strong.

A BFT $\not\Rightarrow$ IC

We present psuedocode for D-S TRB [7] as described in Section 4 in Figure 2. Figure 3 displays pseudocode for the lazy protocol described in Section 4.

In the following Lemmas, we show that n players can still achieve $\mathcal{F}_{TRB}$ even when one player follows λ_r while all remaining players follow $\vec{\delta}$. Similar proofs for $\vec{\delta}$ appear in [7]. As a notational aid, we say that a player p *extracts* m when p inserts m into the set *extracted*. According to the pseudocode for Dolev-Strong and the lazy strategy, non-Byzantine players only extract valid messages.

Lemma 12 (Lazy Agreement)**.** *Suppose non-Byzantine players follow strategy profile* $(\vec{\delta}_{\mathcal{N}-\mathcal{T}-\{r\}}, \lambda_r)$ *and that* $|\mathcal{T}| \leq f$*. If a non-Byzantine processes delivers* m*, then all non-Byzantine processes eventually deliver* m*.*

Proof. We only consider the case where $f > 0$ and $n > f + 2$ where λ_r and δ_r differ.

Let i be the earliest round in which some non-Byzantine process q extracts m in instance k. Since there are at most f Byzantine nodes and a message is extracted only if it is valid, $i \leq f$. We show that all non-Byzantine players either extract m by round $f + 1$ or extract at least 2 distinct values $m' \neq m''$ such that $m \notin \{m', m''\}$.

Consider $i = 0$, implying $leader^k = q$. If q is following λ_q then q sends the value message to $f+1$ nodes in the first round, one of which, say node c, must be non-Byzantine and following δ_c since $n > f + 2$. In round 2, c forwards the message to all other players and all non-Byzantine players extract m in round $2 \leq f + 1$ since $f > 0$, unless they have

already extracted two values. If q instead follows δ_q, then q sends the message to $n-1$ other nodes in the first round and all non-Byzantine players extract the value m in round $1 \leq j < f+1$ or two distinct values m', m'' by round j.

Consider $0 < i \leq f$, implying $\textit{leader}^k \neq q$. If q follows δ_q or $i = f$ (implying $\delta_q = \lambda_q$ in this round), then q forwards the message to $n-f-1$ other nodes in the following round, $i \leq f+1$, and we reach our conclusion. Otherwise, q follows λ_q and $i < f$. In this case, q forwards the message to $f+1-s$ players in round $i+1 \leq f$ where $s \geq i$ since the extracted message contains at least i signatures. Round i is the first round in which a non-Byzantine player extracts m so that only $f-i$ of these players may be Byzantine. Hence, at least one player c to which q forwards the message is non-Byzantine. Player c must follow δ_c because we assumed only q follows λ_q. In the following round $i+2 \leq f+1$, c sends m to all remaining non-Byzantine nodes, and they either extract m in round $i+2$ or have extracted disinct two values $m' \neq m''$ by round $i+2$.

If any non-Byzantine player extracts m, then all non-Byzantine players extract m or two distinct values $m' \neq m''$ by round $f+1$. All non-Byzantine players thus have identical extracted sets and deliver the same value at the end of round $f+1$. □

Lemma 13 (Lazy Integrity)**.** *Suppose non-Byzantine players follow strategy profile $(\vec{\delta}_{\mathcal{N}-\mathcal{T}-\{r\}}, \lambda_r)$ and that $|\mathcal{T}| \leq f$. Every non-Byzantine process delivers at most one message, and if it delivers $m \neq$ **SF** then some process must have broadcast m.*

Proof. Both strategies specify a single message is delivered only during round $f+1$. If $m \neq$ **SF**, then m must have been extracted in round $i \leq f+1$. Only valid messages are extracted, and a message is valid only if it contains the signature of the leader. As signatures are unforgeable, the leader must have broadcast the message m. □

Lemma 14 (Lazy Validity)**.** *Suppose non-Byzantine players follow strategy profile $(\vec{\delta}_{\mathcal{N}-\mathcal{T}-\{r\}}, \lambda_r)$ and that $|\mathcal{T}| \leq f$. If the leader is non-Byzantine and broadcasts m, then all non-Byzantine processes eventually deliver m.*

Proof. A valid message requires a signature by the leader. A non-Byzantine leader following either strategy broadcasts a single value m so that there is only one valid value that is signed by m. Hence, for the leader, $\textit{extracted} = \{m\}$, and the leader delivers m in the final round. By Lemma 12 all other non-Byzantine players also deliver m. □

Lemma 15 (Lazy Termination)**.** *Suppose non-Byzantine players follow strategy profile $(\vec{\delta}_{\mathcal{N}-\mathcal{T}-\{r\}}, \lambda_r)$ and that $|\mathcal{T}| \leq f$. Every non-Byzantine process eventually delivers some message.*

Proof. All strategies terminate in round $f+1$ with a value delivered. □

Theorem 2 (Lazy Safety and Liveness (**TRB1-4**)) *For all $\mathcal{T} \subseteq \mathcal{N}$ such that $|\mathcal{T}| \leq f$, for all $\vec{\tau}_{\mathcal{T}} \in \mathcal{S}_{\mathcal{T}}$, if $\vec{\sigma} = (\vec{\delta}_{\mathcal{N}-\mathcal{T}-\{r\}}, \lambda_i, \vec{\tau}_{\mathcal{T}})$ is played for Γ_{TRB} then* **TRB1-4** *are fulfilled.*

Proof. Follows from Lemmas 13, 14, 12, and 15. □

The proof that D-S TRB is not a Nash equilibrium relies on the observation that a single node following the lazy protocol does not violate safety (and thus does not adversely impact benefit) while requiring the lazy node to pay fewer costs in some instances.

Theorem 3 (Not a Nash Equilibrium) *Consider $\bar{u}_r$ defined for a risk averse rational player r: $\bar{u}_r(\vec{\delta}_{\mathcal{N}-\mathcal{T},\{r\}}, \lambda_r) > \bar{u}_r(\vec{\delta})$ if $n > f+2$.*

Proof. It follows from Theorem 2 that **TRB1-4** are all maintained and thus $\textit{benefits}_r^k(\vec{\delta}_{\mathcal{N}-\{r\}}, \lambda_r) = \textit{benefits}_r^k(\vec{\delta})$ for all $|\mathcal{T}| \leq f$ and for all $\vec{\tau}_{\mathcal{T}} \in \mathcal{S}_{\mathcal{T}}$.

Worst-case Byzantine behavior for a specified strategy $\vec{\sigma} \in \mathcal{S}_{\mathcal{N}}$ minimizes utility as defined by the risk averse $\bar{u}_r$ by maximizing $\textit{costs}_r^k(\vec{\sigma})$, which depends upon $\mathrm{sent}_r^k(\vec{\sigma})$. We show that $\mathrm{sent}_r^k(\vec{\delta}_{\mathcal{N}-\{r\}}, \lambda_r) \subseteq \mathrm{sent}_r^k(\vec{\delta})$ for all $k > 0, f \geq 0, |\mathcal{T}| \leq f$, and $\vec{\tau}_{\mathcal{T}} \in \mathcal{S}_{\mathcal{T}}$. If $\textit{leader}^k = r$, then both

protocols specify the same message m, but λ_r specifies fewer when $n > f + 2$ by sending m to $f + 1$ players while δ_r specifies $n - 1$. If *leader*$^k \neq r$, then $\text{sent}^{\vec{\sigma}}_{r\to k}(f)$or some profile $\vec{\sigma} \in \mathcal{S}_\mathcal{N}$ depends upon the set of messages received since messages are never replied to (that is, a message with value v is not sent to members of sig_v). As all other players follow the same profile $\vec{\delta}_{\mathcal{N}-\{r\}}$ in either case, player r receives the same set of messages in either $\vec{\delta}$ or $(\vec{\delta}_{\mathcal{N}-\{r\}}, \lambda_r)$. The set of messages sent in each profile is identical though λ_r specifies fewer messages containing each value to be sent. Indeed, λ_r specifies $\min(n-1, f+1)-s$ messages while δ_r specifies $n-1-s$ messages, implying $\text{sent}^k_r(\vec{\delta}_{\mathcal{N}-\{r\}}, \lambda_r) \subset \text{sent}^k_r(\vec{\delta})$ when $n > f + 2$. Hence, when $n > f + 2$, $\textit{costs}^k_r(\vec{\delta}_{\mathcal{N}-\{r\}}, \lambda_r) < \textit{benefits}^k_r(\vec{\delta})$.

□

The proof that the lazy protocol is not a viable alternative to D-S TRB if all players follow it is based on exhibiting an execution in which **TRB3** is violated.

Theorem 4 (Failed Agreement) *If non-Byzantine players follow strategy profile $\vec{\lambda}_{\mathcal{N}-\mathcal{T}}$, $|\mathcal{T}| \leq f$, $f > 1$ and $n > f + 2$ then* **TRB3** *can fail in all instances of* Γ_{TRB}.

Proof. Suppose a non-Byzantine player r is the leader. In the first round, r sends the broadcast value $\langle \text{VALUE}, v, k\rangle_r$ to $f + 1$ other players. WLOG, assume f of these players are Byzantine players that never forward f in order to reduce u^k_r, and the remaining player p is non-Byzantine. Since $f > 1$ there are guaranteed to be at least three rounds. In the second round, p also sends $\langle \text{VALUE}, v, k\rangle_{r,p}$ to $f + 1$ players. Because $sigs_m = \{r\}$ for this message, p sends the message to only f other players, which may be the same set of Byzantine players chosen by r. Thus, the value is again not forwarded to any other players. Since $n > f + 2$ there exists a third non-Byzantine player g that never receives m while r and p both deliver m in the final round. □

We now provide the complete proof that D-S TRB is (k, t)-robust. if $c_{snd}(m) = 0$.

Theorem 5 (Zero Cost D-S TRB) *D-S TRB is (k, t)-robust if $c_{snd}(m) = 0$ for all messages m.*

Proof. The function $u^j_i(\vec{\delta}_{\mathcal{N}-\mathcal{T}}, \vec{\tau}_\mathcal{T})$ for any instance j of Γ_{TRB} is defined as $\textit{benefits}^j_i(\vec{\delta}_{\mathcal{N}-\mathcal{T}}, \vec{\tau}_\mathcal{T}) - \textit{costs}^j_i(\vec{\delta}_{\mathcal{N}-\mathcal{T}}, \vec{\tau}_\mathcal{T})$. By assumption $\textit{costs}^j_i(\vec{\delta}_{\mathcal{N}-\mathcal{T}}, \vec{\tau}_\mathcal{T})=0$ for all $j>0$. In [7], Dolev and Strong prove that $\vec{\delta}_{\mathcal{N}-\mathcal{T}}$ fulfills $\mathcal{F}_{TRB}$ by meeting **TRB1–4** for all $|\mathcal{T}| \leq t$ and for all $\vec{\tau}_\mathcal{T} \in \mathcal{S}_\mathcal{T}$, implying for any strategy profile $\vec{\sigma} \in \mathcal{S}_\mathcal{N}$ for Γ_{TRB}, $\textit{benefits}^j_i(\vec{\sigma}_{\mathcal{N}-\mathcal{T}}, \vec{\tau}_\mathcal{T}) \leq \textit{benefits}^j_i(\vec{\delta}_{\mathcal{N}-\mathcal{T}}, \vec{\tau}_\mathcal{T})$. Hence, any coalition strategy $\vec{\phi}_\mathcal{C} \in \mathcal{S}_\mathcal{C}$ for $\mathcal{C} \subseteq \mathcal{N}$ such that $\mathcal{C} \cap \mathcal{T} = \emptyset$ and $|\mathcal{C}| \leq k$ must result in at most utility $u^j_i(\vec{\delta}_{\mathcal{N}-\mathcal{T}}, \vec{\tau}_\mathcal{T})$ implying that $u^j_i(\vec{\delta}_{\mathcal{N}-\mathcal{T}}, \vec{\tau}_\mathcal{T}) \geq u^j_i(\vec{\delta}_{\mathcal{N}-(\mathcal{C}\cup\mathcal{T})}, \vec{\phi}_\mathcal{C}, \vec{\tau}_\mathcal{T})$ for all $j>0$. □

B BaN TRB

BaN TRB is very closely related to D-S TRB. The close relationship between the two protocols leads to the following Lemma. The intuition for the Lemma is that the same set of *VAL* messages are forwarded by a collection of players playing $\vec{\rho}$as would be forwarded by players plaing $\vec{\delta}$.

Lemma 16. *If all non-Byzantine players follow the BaN TRB protocol, then they deliver the same values that they would deliver if they ran D-S TRB instead.*

Proof. It follows from Lemma 1 that every pair of non-Byzantine players are friends.

Consider the case with a non-Byzantine leader. The leader sends exactly one value v to all non-Byzantine players in round 1 of both protocols, so every non-Byzantine player receives v in both protocols and delivers it in round $f + 1$.

With a Byzantine leader, if any non-Byzantine player delivers v broadcast by the leader, then some non-Byzantine player received any value broadcast by the sender by round f at the latest (since there are at most f Byzantine players). Both protocols specify that any non-Byzantine player forward the first 2 values it receives in an instance to all other players. So at the end of round $f + 1$ all players will have received either the same unique value v,

or at least two values, or no value. Non-Byzantine players all deliver **SF** in the latter two cases or the unique value v in the first case. □

Theorem 6 *For all $\mathcal{T} \subseteq \mathcal{N}$ such that $|\mathcal{T}| \leq f$, for all $\vec{\tau}_{\mathcal{T}} \in \mathcal{S}_{\mathcal{T}}$, if $\vec{\sigma} = (\vec{\rho}_{\mathcal{N}-\mathcal{T}}, \vec{\tau}_{\mathcal{T}})$, then $\mathbf{TRB1-4}$ hold in all instances of Γ_{TRB} when $\vec{\sigma}$ is played.*

Proof. Follows from Lemma 16 and D-S TRB maintains $\mathbf{TRB1-4}$. □

C BaN TRB Analysis

In Section 6 we defined friends, ex-friends, and enemies. Here we prove properties relating these sets to strategy profile $\vec{\sigma}$.

Lemma 1 *Suppose two players p and q play ρ_p and ρ_q, respectively. For all $\vec{v}_{\mathcal{N}-\{p,q\}} \in \mathcal{S}_{\mathcal{N}-\{p,q\}}$, let $\vec{\sigma} = (\vec{\rho}_{\{p,q\}}, \vec{v}_{\mathcal{N}-\{p,q\}})$. $p \in \mathrm{F}_q(\vec{\sigma})$.*

Proof. By definition, for all $k > 0$, $\mathrm{seq}^k_{p\to q}(\vec{\sigma}) \in \mathcal{M}^k_{p\to q}$, implying $p \in \mathrm{F}_q(\vec{\sigma})$. □

Lemma 2 *Suppose player p plays ρ_p and for all $\vec{v}_{\mathcal{N}-\{p\}} \in \mathcal{S}_{\mathcal{N}-\{p\}}$, let $\vec{\sigma} = (\rho_p, \vec{v}_{\mathcal{N}-\{p\}})$. If $q \in \mathcal{N}$ and $p \notin \mathrm{F}_q(\vec{\sigma})$, then $p \in \mathrm{E}_q(\vec{\sigma})$.*

Proof. Since p follows ρ_p, p by definition sends an acceptable message sequence to players not in shun_p. By the assumption that $p \notin \mathrm{F}_q(\vec{\sigma})$, we infer that $q \in \mathit{shun}_p$, and thus p does not send a VAL_1 message to q when p is leader. Hence, $p \in \mathrm{E}_q(\vec{\sigma})$. □

C.1 Faithful Utility

We address the corner cases of $n = f + 1, f > 0$ and $f = 0$ that we omitted in Section 6. The proofs for the former case are similar in flavor to those for the case when $n > f + 1, f > 0$, differing only in the fact that $|\mathcal{T}| = f - 1$ rather than f.

C.1.1 $f > 0$ and $n = f + 1$

In the first corner case in which $n = f + 1$ and $f > 0$, r pays a steady state cost of at most $\mathrm{C}(n-f, f-1)$ and $\bar{u}_i r\vec{\rho} = \frac{(\beta+n\varpi)-\mathrm{C}(n-f-1,f)}{n}$.

Lemma 17. *Let $n = f + 1$ and $f > 0$. For all $\mathcal{T} \subseteq \mathcal{N}$, $|\mathcal{T}| \leq f$ and $\forall \vec{\tau}_{\mathcal{T}} \in \mathcal{S}_{\mathcal{T}}$, let $\vec{\varphi} = (\vec{\rho}_{\mathcal{N}-\mathcal{T}}, \vec{\tau}_{\mathcal{T}})$ and rational player $r \notin \mathcal{T}$. $\mathrm{costs}_r(\vec{\varphi}) \leq \mathrm{C}(n - f, f - 1)$*

Proof. It follows from the definition of $\vec{\varphi}$ and Lemma 1 $|\mathrm{F}_r(\vec{\varphi})| \geq n - f$ and $|\mathrm{E}_r(\vec{\varphi})| \leq f - 1$. It follows from Lemma 4 that $\mathrm{C}(x, y)$ is maximized when $x = f - 1$. Since $\mathrm{C}(x, y)$ is defined for ρ, $\mathit{costs}_r(\vec{\varphi}) \leq \mathrm{C}(n - f, f - 1)$. □

Ultimately, the Byzantine aware utility for the case when $n = f + 1$ and $f > 0$ is $\frac{(\beta+n\varpi)-\mathrm{C}(n-f,f-1)}{n}$.

Lemma 18. *Let $n = f + 1$ and $f > 0$. For all $r \in \mathcal{N}$, $\bar{u}_r(\vec{\rho}) = \frac{(\beta+n\varpi)-\mathrm{C}(n-f,f-1)}{n}$*

Proof. It follows from Lemma 17 that $\mathit{costs}_r(\vec{\varphi}) \leq \mathrm{C}(n - f, f - 1)$ for $\vec{\varphi} = (\vec{\rho}_{\mathcal{N}-\mathcal{T}}, \vec{\tau}_{\mathcal{T}})$ for all $\mathcal{T} \subseteq \mathcal{N} - \{r\}$ such that $|\mathcal{T}| \leq f - 1$, for all $\vec{\tau}_{\mathcal{T}} \subseteq \mathcal{S}_{\mathcal{T}}$.

It follows from Lemma 3 that the steady state utility for r is at most $\beta + n\varpi - \mathrm{C}(n - f, f - 1)$. Division by n completes the proof. □

C.1.2 $f = 0$

The special case when $f = 0$ is slightly more interesting as it does not follow the same structure as the two $f > 0$ cases. One primary difference is the cost required to maintain friends changes since only the leader is required to send messages.

Lemma 19. *Let $f = 0$. For all $r \in \mathcal{N}$, $\mathrm{costs}_r(\vec{\rho}) = (n - 1)\gamma$.*

Proof. It follows from Lemma 1 that $\mathcal{N} - \{r\} = \mathrm{F}_r(\vec{\rho})$. The protocol requires a single message to each other participant only during instances when r is leader, so the cost to r is $(n - 1)\gamma$. □

The Byzantine aware utility for following the profile when $f = 0$ is $\frac{(\beta+n\varpi)-(n-1))\gamma}{n}$.

Lemma 20. *Let $f = 0$. For all $r \in \mathcal{N}$, $\bar{u}_r(\vec{\rho}) \geq \frac{(\beta+n\varpi)-(n-1))\gamma}{n}$.*

Proof. Since $f = 0$ the only strategy profile to be

considered is $\vec{\rho}$. It follows from Lemma 19 that $costs_r(\vec{\rho}) \leq (n-1)\gamma$. It follows form Lemma 3 that the steady state utility for r is at most $\beta + n\varpi - \mathrm{C}(n-f, f-1)$. Division by n completes the proof. □

C.2 Deviant Utility

First, we dispatch a Lemma that was only sketched in the main sections. The following Lemma shows that a rational player with no friends cannot obtain more than ϖ benefit.

Lemma 9 *Let $n \geq f+1$, $f > 0$. For all $\mathcal{T} \subseteq \mathcal{N}$, $|\mathcal{T}| \leq f$ and $\forall \sigma_r \in \mathcal{S}r$, $\forall r \in \mathcal{N} - \mathcal{T}$, let $\vec{o} = (\sigma_r, \vec{\rho}_{\mathcal{N}-\mathcal{T}-\{r\}}, \vec{\varsigma}^r_{\mathcal{T}})$. If $|\mathrm{F}_r(\vec{o})| = 0$ and $|\mathrm{X}^{\mathrm{NON}}_r(\vec{o})| > 0$, then* benefits$_r(\vec{o}) \leq \varpi$

Proof. The spiteful profile $\vec{\varsigma}^r_{\mathcal{T}}$ is similar to $\vec{\rho}_{\mathcal{T}}$, but shuns r. Further, by Lemma 2, $\mathrm{X}^{\mathrm{NON}}_r(\vec{o}) \subseteq \mathrm{E}_r(\vec{o})$, implying $\forall q \in \mathrm{X}^{\mathrm{NON}}_r(\vec{o}) : r \in shun_q$. Hence, all players besides r send a single value to all other players except r. The lower bound on $\mathrm{X}^{\mathrm{NON}}_r(\vec{o})$ implies at least one other non-Byzantine player q receives the single value in each instance from all leaders other than r and delivers the value in the last round. Player r instead receives no messages and thus delivers **SF**, violating **TRB3** when r is not leader, implying no benefit is obtained for these instances. When r is leader, it sends no messages. All other players, receiving no messages, deliver **SF**, which violates **TRB3** unless r also delivers **SF**. Since r can never deliver a value $\neq$ **SF** and guarantee **TRB3**, the β benefit cannot be obtained so that $benefits_r(\vec{o}) \leq \varpi$. □

Addressing the deviations is more complicated in the corner cases as there are multiple types of deviations that a rational player may try in order to increase cost.

C.2.1 $n = f+1$ and $f > 0$

When $n = f+1$, the proof structure is identical to the case when $n > f+1$ except that the optimal number of Byzantine players is $f-1$ rather than f.

Lemma 21. *Let $f > 0$, $n = f+1$. For all $\mathcal{T} \subseteq \mathcal{N}$, $|\mathcal{T}| = f-1$ and $\forall r \notin \mathcal{T}$, $\forall \sigma_r \in \mathcal{S}r$, let $\vec{o} = (\sigma_r, \vec{\rho}_{\mathcal{N}-\mathcal{T}-\{r\}}, \vec{\varsigma}^r_{\mathcal{T}})$. If $|\mathrm{E}_r(\vec{o})| < n-1$, then* costs$_r(\vec{o}) \geq \mathrm{C}(n-f, f-1)$.

Proof. It follows from Lemma 7 that $\mathcal{T} \subseteq \mathrm{E}_r(\vec{o})$ and from Lemma 2 that $\forall p \in \mathcal{N}, \notin \mathcal{T} - \{r\} \implies p \notin \mathrm{X}_r(\vec{o}) - \mathrm{E}_r(\vec{o})$. Hence $|\mathrm{E}_r(\vec{o})| \geq f$, $|\mathrm{F}_r(\vec{o})| \leq n-f$, and $|\mathrm{F}_r(\vec{o})| + |\mathrm{E}_r(\vec{o})| = n-f-1$. It thus follows from Lemma 4 that $\mathrm{C}(|\mathrm{F}_r(\vec{o})|, |\mathrm{E}_r(\vec{o})|)$ is minimized when $|\mathrm{F}_r(\vec{o})| = n-f$. Thus, $costs_r(\vec{o}) \geq \mathrm{C}(n-f, f-1)$. □

Lemma 22. *Let $f > 0$, and $n = f+1$. For all $r \in \mathcal{N}$, $\bar{u}_r(\vec{\rho}_{\mathcal{N}-\{r\}}, \sigma_r) \leq \max\{\frac{(\beta+n\varpi)-\mathrm{C}(n-f,f-1)}{n}, \frac{\varpi}{n}\}$*

Proof. There exists strategy profile $\vec{o} = (\vec{\rho}_{\mathcal{N}-\mathcal{T}-\{r\}}, \sigma_r, \vec{\varsigma}^r_{\mathcal{T}})$ such that $\mathcal{T} \subseteq \mathcal{N} - \{r\}$ and $|\mathcal{T}| = f-1$.

Consider the case where $|\mathrm{E}_r(\vec{o})| < n-1$. It follows from Lemma 21 part 1 that $costs_r(\vec{o}) \geq \mathrm{C}(n-f, f-1)$. It follows from the definition of benefits in Γ_{TRB} that the steady state benefits for r of $\vec{o}$ are at most $\beta + n\varpi$. Hence $\bar{u}_r(\vec{o}) \leq \frac{(\beta+n\varpi)-\mathrm{C}(n-f,f-1)}{n}$.

Consider the case where $|\mathrm{E}_r(\vec{o})| = n-1$. It follows Lemma 9 that $benefits_r(\vec{o}) \leq \varpi$. Hence $\bar{u}_r(\vec{o}) \leq \frac{\varpi}{n}$.

Combining the two cases, $\bar{u}_r(\vec{\rho}_{\mathcal{N}-\{r\}}, \sigma_r) \leq \max\{\frac{(\beta+n\varpi)-\mathrm{C}(n-f,f-1)}{n}, \varpi\}$. □

C.2.2 $f = 0$

The cost a deviant player must pay is proportional to the number of friends the player maintains.

Lemma 23. *Let $f = 0$. For all $r \in \mathcal{N}$, $\forall \sigma_r \in \mathcal{S}r$, let $\vec{o} = (\vec{\rho}_{\mathcal{N}-\{r\}}, \sigma_r)$.* costs$_r(\vec{o}) \geq |\mathrm{F}_r(\vec{o})|\gamma$

Proof. It follows from Lemma 1 that r must send one *VAL* message to each $p \in \mathrm{F}_r(\vec{o})$ during each instance that r is leader in order to maintain the friendset, incurring cost at least $|\mathrm{F}_r(\vec{o})|\gamma$. □

Now we can figure out the aware utility.

Lemma 24. *Let $f = 0$. For all $r \in \mathcal{N}$, let $\vec{o} = (\vec{\rho}_{\mathcal{N}-\{r\}}, \sigma_r)$. $\bar{u}_r(\vec{o}) \leq \max\{\frac{\beta+n\varpi-(n-1)\gamma}{n}, \frac{\varpi}{n}\}$*

Proof. Consider the case where $|\mathrm{E}_r(\vec{o})| < n-1$. It follows from $\beta + n\varpi$ being the maximal benefit achievable in the game and from Lemma 23 that $\bar{u}_r(\vec{o}) \leq \beta + n\varpi - (n-1)\gamma$.

Consider the case where $|\mathrm{E}_r(\vec{o})| = 0$. It follows from Lemma 9 that $\mathit{benefits}_r(\vec{o}) \leq \varpi$ and consequently that $\bar{u}_r(\vec{o}) \leq learn$. □

And finally we can show that the protocol is a Nash equilibrium.

C.2.3 $n = 1$

For completeness we observe that $n = 1$ is trivial and non-interesting.

C.3 Nash Equilibria

We can finally show that BaN TRB is a Nash Equilibrium when $n = f + 1$ and $f > 0$.

Theorem 9. *Using the risk-averse rational model, $\vec{\rho}$ is a Nash equilibrium if $\beta + (n-1)\varpi \geq \mathrm{C}(n-f, f-1)$, $n = f+1$ and $f \geq 0$.*

Proof. It suffices to show $\forall i \in \mathcal{N}, \forall \sigma_i \in \mathcal{S}i$ $\bar{u}_i(\vec{\rho}) \geq \bar{u}_i(\vec{\rho}_{\mathcal{N}-\{i\}}, \sigma_i)$.

It follows from 18 that $\bar{u}_i(\vec{\rho}) \geq \frac{\beta+n\varpi-\mathrm{C}(n-f,n-f-1)}{n}$ and from 22 that $\bar{u}_i(\vec{\rho}_{\mathcal{N}-\{i\}}, \sigma_i) \leq \max\{\frac{\beta+n\varpi-\mathrm{C}(n-f,f-1)}{n}, \frac{\varpi}{n}\}$. It thus follows from our assumption that $\beta + (n-1)\varpi \geq \mathrm{C}(n-f, f-1)$ that $\bar{u}_i(\vec{\rho}) \geq \bar{u}_i(\vec{\rho}_{\mathcal{N}-\{r\}}, \sigma_i)$, completing the proof. □

And for $f = 0$.

Theorem 10. *Let $f = 0$ and $n > 1$. Using the risk-averse rational model, $\vec{\rho}$ is a Nash equilibrium if $\beta + (n-1)\varpi \geq (n-1)\gamma$.*

Proof. It suffices to show $\forall i \in \mathcal{N}, \forall \sigma_i \in \mathcal{S}i$ $\bar{u}_i(\vec{\rho}) \geq \bar{u}_i(\vec{\rho}_{\mathcal{N}-\{i\}}, \sigma_i)$.

It follows from 18 that $\bar{u}_i(\vec{\rho}) \geq \max\{\frac{\beta+n\varpi-(n-1)\gamma}{n}, \frac{\varpi}{n}\}$ and from 24 that $\bar{u}_i(\vec{\rho}_{\mathcal{N}-\{i\}}, \sigma_i) \leq \max\{\frac{\beta+n\varpi-(n-1)\gamma}{n}, \frac{\varpi}{n}\}$. It thus follows from our assumption that $\beta + (n-1)\varpi \geq (n-1)\gamma$ that $\bar{u}_i(\vec{\rho}) \geq \bar{u}_i(\vec{\rho}_{\mathcal{N}-\{r\}}, \sigma_i)$, completing the proof. □

Cost of Byzantine Anarchy and Malice. When $n = f+1$ and $f > 0$, $PoB(f) = \frac{\beta+n\varpi-\mathrm{C}(n-f,f-1)}{\beta+n\varpi-(n-1)\gamma}$ and $PoM(f) = \frac{\beta+n\varpi-\mathrm{C}(n-f,f-1)}{\beta+n\varpi-2n(n-1)\gamma}$. When $f = 0$, $PoB(f) = PoM(f) = 1$, both by definition. In general, the Price of Malice is $\bar{u}_i r[\vec{\rho}]$ divided by the optimal cost of the trivial 1 round 0 fault tolerant TRB protocol. This reflects the additional cost rational players have to pay when faulty players are allowed into the system.. While the Price of Byzantine Anarchy