

Rationality in the Full-Information Model

Ronen Gradwohl*

Kellogg School of Management, Northwestern University, Evanston, IL 60208
r-gradwohl@kellogg.northwestern.edu

Abstract. We study rationality in protocol design for the full-information model, a model characterized by computationally unbounded adversaries, no private communication, and no simultaneity within rounds. Assuming that players derive some utility from the outcomes of an interaction, we wish to design protocols that are faithful: following the protocol should be an optimal strategy for every player, for various definitions of “optimal” and under various assumptions about the behavior of others and the presence, size, and incentives of coalitions. We first focus on leader election for players who only care about whether or not they are elected. We seek protocols that are both faithful and resilient, and for some notions of faithfulness we provide protocols, whereas for others we prove impossibility results. We then proceed to random sampling, in which the aim is for the players to jointly sample from a set of m items with a distribution that is a function of players’ preferences over them. We construct protocols for $m \geq 3$ that are faithful and resilient when players are single-minded. We also show that there are no such protocols for 2 items or for complex preferences.

1 Introduction

The full-information model of Ben-Or and Linial [8] is one of the classically-studied settings for protocol design. In this model there are no computational limits on the adversary, there is no private communication, and there is no guarantee of simultaneity within rounds of a protocol. Three famous problems are collective coin-flipping, leader election, and random sampling. In the first, players jointly flip a coin; in the second, they jointly select a random player; and in the third, they jointly select a random element from some universe of m items. In general, the goal is to design protocols that are resilient: the outcome should be random even in the presence of an adversary who corrupts and coordinates the behavior of a fraction of the players.

In this paper we explore the role of preferences in the design of such protocols. While preferences are not explicitly considered in the well-studied formulations of the problems, they are implicitly present. For example, leader election has a fairness criterion, which requires each player to be elected with roughly equal

The original version of this chapter was revised: The copyright line was incorrect. This has been corrected. The Erratum to this chapter is available at DOI: [10.1007/978-3-642-11799-2_36](https://doi.org/10.1007/978-3-642-11799-2_36)

* Much of this research was done while the author was a graduate student at the Weizmann Institute of Science, supported by US-Israel Binational Science Foundation Grant 2002246 and ISF Grant 334/08.

probability (presumably because everybody wants to be the leader). A leader election protocol is resilient if an adversary can not force the elected leader to be a member of his coalition (or at least will fail to do so with constant probability). Again, the adversary *wants* a coalition-member to be elected. For collective coin-flipping and random sampling, resilience is also measured as a bound on the probability that an adversary succeeds at something. It is implicitly assumed that the adversary wants to do this, and that the honest (non-adversarial) players do not wish him to achieve his goal.

The study of preferences in the design of protocols is primarily the domain of mechanism design. In mechanism design a planner wishes to implement some function of players' private information. His goal is to design a mechanism and provide incentives for the players so that their optimal strategy is to truthfully reveal their private information, and more generally to adhere to the mechanism. The optimality of players' strategies is measured via some solution concept: following the mechanism should be in some equilibrium, most commonly Nash, ex post Nash, or dominant strategy. In this paper we take a similar approach – we define new solution concepts appropriate for the full-information model, and seek protocols that are *faithful*: following them is optimal for players with respect to these solution concepts (in addition to the usual resilience guarantees).

For any problem of protocol design, making the structure of preferences explicit has two potential benefits, both of which we achieve in this paper. First, it can result in better protocols – protocols are arguably of little use if players have no incentives to follow them. If one can obtain faithful protocols without harming the original guarantees of the protocol, then one has only gained. Second, it may be possible to sidestep some impossibility results of the original problem, since often these impossibility results are based on arbitrary play by the adversary. If players do not play arbitrarily but rather obey some preference structure, then many of these results no longer hold.

The model. In the full-information model all communication is by broadcast. In each round, some of the players send a message, which may depend on messages sent in previous rounds. The main difficulty is that adversarial players are allowed to “rush” – to wait until all messages have been sent within a round, and only then to send their own messages.

This paper. We are largely motivated by recent work in rational cryptography, in which the aim is to design cryptographic protocols that participants *want* to follow. Two of the main difficulties encountered when attempting a game theoretic analysis of cryptographic protocols are computational limits and potentially adversarial timing. In this paper we focus solely on the latter issue by considering a model in which (adversarial) players may be computationally unbounded, and the guaranteed security (i.e. resilience) is information-theoretic. We highlight the various challenges and subtleties caused by a combination of rational and adversarial players, particularly in the presence of adversarial timing. We also draw a possibility-impossibility border for various problems and requirements in this setting. Finally, we believe that this paper is an illuminating stepping-stone towards a game theoretic analysis of more general cryptographic protocols.

1.1 Our Results and Organization

Definitions (Section 2). The initial difficulty encountered when considering preferences in the full-information model is to precisely formulate a notion of equilibrium. The first notion to consider is Nash equilibrium (NE), in which each player's strategy must be expected utility maximizing assuming others also follow their Nash strategies. If the protocol is such that only one player sends a message in each round, then this suffices. One such protocol is Baton Passing [32], a protocol that is resilient and in fact satisfies our weaker solution concept¹. However, state-of-the-art protocols are often round-efficient, and allow multiple players to broadcast within a round. Because of the lack of synchrony within rounds, however, NE does not suffice. In the Lightest Bin protocol [14], for example, a player may increase his chance of winning from $1/n$ (where n is the number of players) to a constant by deviating. To deal with asynchrony, we will require that for *any ordering* of the players within each round, the protocol is in a NE. In Section 2 we formalize this and other notions of what it means to be faithful and faithful in the presence of adversarial players.

Impossibility with complex preferences (Section 3). In Section 3 we encounter our first impossibility result. Theorem 1 states that no random selection protocol can satisfy even our weakest solution concept if players have a full preference order over the outcomes of the protocol. One implication of this impossibility result is that collective coin-flipping is impossible with players who have some preference about the outcome. For leader election and random sampling, this result forces us to concentrate on more restricted preferences for players. For the former, we assume that players care only about whether or not they are elected, and are indifferent otherwise. For the latter, we assume players are single-minded: each prefers one of the items, and is indifferent about the others.

Faithfulness with resilience (Section 4.2). The standard aim of selection protocols in the full-information model is resilience: if an adversary corrupts and coordinates the actions of a fraction of the players, he still fails to force his desired outcome with non-negligible probability. In Section 4.2 we construct optimal protocols that both satisfy a notion of equilibrium and are resilient. Players wish to faithfully adhere to the protocol if the others also do, and there is a resilience guarantee in the presence of an adversary.

Faithfulness in the face of an adversary (Section 4.3). In Section 4.3 we consider the problem of constructing leader election protocols that are in equilibrium even when not all others follow the protocol. We show that it is impossible to construct such protocols in the presence of a malicious adversary, even if the adversary has his own objective of maximizing the probability that a coalition-member wins. However, if the adversary maximizes this probability, but also only deviates from the protocol if he strictly gains from doing so (i.e. if deviating is costly), then we do design a resilient protocol.

¹ More specifically, it is in a full-information ex post NE – see Definition 5. It is not, however, in a full-information dominant strategy equilibrium (Definition 3).

Resilience to rational coalitions (Section 4.4). A different form of resilience against adversarial play is when there is no controlling adversary, but instead players may form “rational coalitions” to benefit all members. In Section 4.4 we give an impossibility result for one notion of a “rational coalition”, but for a weaker notion provide a protocol that is resilient against all such coalitions of size at most $n - 2$ [2].

Random sampling (Section 5). Our final set of results concerns random sampling. Each player has some preferences over a universe of m items, and the goal is to design a protocol in which an item is sampled with a probability distribution that is a function of those preferences. We design protocols that are simultaneously in a full-information ex post Nash equilibrium (in which truthful revelation of one’s preferences is optimal) and resilient against adversarial coalitions.

1.2 Related Work

This paper draws from three different literatures – protocol design in cryptography and distributed computing, and algorithmic mechanism design. The extensive literature on collective coin-flipping, random sampling, and leader election in the full-information model includes [32, 16, 27, 3, 11, 12, 17, 31, 14, 13, 33, 6]. The paper most closely related to ours is that of Antonakopoulos [6], who also considers 1-round protocols in which individual players have no incentive to deviate. However, his protocols all attain either faithfulness or resilience, but never both. Similarly, Ben-Or and Linial [8] have a 2-round protocol that is faithful but not resilient to larger coalitions. The paper most closely related to ours from the mechanism design literature is that of Altman and Tennenholtz [5], who construct 1-round protocols that are faithful (but also not resilient). Their goal is to attain arbitrary distributions over the players. Also related is the literature on ranking games [10], in which players have preferences about their rankings in some game.

While we believe that we are the first to study notions of rationality tailored specifically for the full-information model, such notions have been studied in other settings for distributed computing. For example, Monderer and Tennenholtz [25] consider an implementation problem in a distributed network. Shneidman and Parkes [34] introduce the idea that protocols should be faithful. Additionally, the field of Distributed Algorithmic Mechanism Design (DAMD) focuses on implementing mechanisms for various problems in a distributed setting. In a general “mission statement” for DAMD, Feigenbaum and Shenker [15] argue that it would be desirable to incorporate various fault models into the DAMD framework. Also, Halpern [21, 22] has expressed the need to incorporate faulty or malicious behavior into distributed settings with rational players. Some papers that address this issue are Aiyer et al. [2], Abraham et al. [1], and Gradwohl [18].

Finally, as mentioned in the introduction, this work is closely related to the growing literature on rational cryptography (see, for example, Katz [24] and

² Compare this with the fact that there are *no* protocols that are resilient against an adversary of size $n/2$ [32].

the references therein). Many works in this literature study rational behavior in a cryptographic setting, for which the full-information model is a special case. However, due to computational issues, the definitions in the general setting are messier (and often also weaker). We note that the way we model rushing is closely related to an idea of Ong et al. [29], who adopt the methods of Kalai [23] to a protocol design setting. The idea of considering rational coalitions was also explored in this context by Ong et al. [28].

Our notions of stability of coalitions are related to similar notions in the game theory literature, such as the strong Nash equilibrium of Aumann [7] and the coalition-proof equilibrium notions of Bernheim et al. [9], Moreno and Wooders [26], and Abraham et al. [1].

2 Protocols and Games

For any vector $X = (X_1, \dots, X_n)$ and $S \subset [n]$, we denote by $X_S = \{X_i\}_{i \in S}$ and by $X_{-S} = \{X_i\}_{i \notin S}$.

2.1 Resilient Protocols

We are interested in protocols involving many players and the incentives of players in following these protocols. Thus, we will assume that players have preferences over possible outcomes, as well as other private information. As in the game theory literature, all this information is collectively called a player’s type. Player i ’s type is denoted by t_i , and the vector $t = (t_1, \dots, t_n)$ is called the type profile. The space of possible types of player i is T_i .

Definition 1 (selection protocol). *An n -player selection protocol \mathcal{P} is specified by a function f , a natural number q , and, for each of the n players, a set of q randomized functions $\{S_i^1, \dots, S_i^q\}_{i \in [n]}$. The protocol proceeds as follows:*

- At round j , the i ’th player broadcasts a random message M_i^j obtained by applying the randomized function S_i^j to all previous messages sent, namely $\{M_k^l : k \in [n], l < j\}$, as well as player i ’s type t_i . The randomness of the function comes from the player’s independent coins.
- After q rounds, the players output $f(\{M_k^l : k \in [n], l \in [q]\})$ which is an element of $[m]$ in an m -item random sampling protocol and an element of $[n]$ in a leader election protocol. If all players follow the protocol then the output is a uniformly random element (unless stated otherwise).

In any round j , a player i ’s legal messages are those in $\bigcup_{t_i \in T_i} \text{supp}(S_i^j(t_i, \{M_k^l : k \in [n], l < j\}))$ ³. We assume that if a player noticeably deviates from the protocol (by broadcasting a message that is not legal), then his message is changed to some default legal value.

³ Note that a player’s legal actions include messages in the support of S_i^j for all types, not just the true one. This is so because the other players do not know i ’s true type.

Players may not legally base their messages in round j on the messages of other players in round j . However, since we can not guarantee simultaneity within a round, we allow the dishonest players to *rush*: they may base their messages on the messages of other players from the same round (but not from later rounds). A leader election protocol is ε -resilient to coalitions of size t if the following holds: If at most t players are playing a coordinated rushing strategy, then the probability that the elected leader is a cheating player is at most $1 - \varepsilon$. Often we will implicitly be referring to a family of protocols, one for each value of n . In this case, we say a protocol is resilient if there exists some $\varepsilon > 0$ such that all protocols in the family with enough players are ε -resilient.

A protocol is *oblivious* if players' messages are based only on their internal coin tosses. A protocol is *explicit* if players' messages and the function f are computable in probabilistic polynomial time (in the number of players and $\log(m)$).

2.2 Extensive-Form Games and Protocols

An n -player *extensive-form game* is specified by a game tree in which every node is owned by a player and outgoing edges are labelled by actions. The game begins at the root node and proceeds down the tree – at every node following the edge labelled by the action played by the node's owner. Payoffs for players are specified at the leaves.

Definition 2 (Nash equilibrium (NE)). A Nash equilibrium (NE) in an extensive-form game is a mixed strategy for every player at every node that he owns, such that: if all players play their NE strategy, then no player obtains a higher expected payoff by deviating at any of his nodes.

We note that in the games we consider, the NE will be completely mixed strategies (i.e. players will play every action with positive probability). Such Nash equilibria are in fact subgame perfect (see [30]).

Consider a selection protocol, where each player i derives some utility $u_i : T_i \times [m] \mapsto \mathbb{R}$ from outputs of the protocol. u_i is such that for $o \neq o' \in [m]$, we have that $u_i(t_i, o) > u_i(t_i, o')$ if and only if player i of type t_i strictly prefers o to o' . Then any protocol in which only one player sends a message in each round can be viewed as an extensive-form game⁴: if after $j - 1$ rounds and messages M_1, \dots, M_{j-1} player i plays in round j , i owns the node at level j in the game tree reached by the game path M_1, \dots, M_{j-1} . Player i 's payoff from an instance of play resulting in o is $u_i(t_i, o)$. Such a selection protocol is in a NE if, in the associated game, it is a NE for every player i to play according to strategy S_i^j if any of his nodes at level j is reached (we say that i follows strategy S_i).

2.3 Rationality in Selection Protocols – Definitions

We now define notions of what it means for a protocol to be faithful, i.e. in which it is in players' best interests to follow the protocol specification. Because

⁴ While it is possible to model simultaneous play as an extensive-form game with *imperfect* information, the ability to rush and the lack of synchrony are more difficult to incorporate into this framework.

there is no synchrony within rounds of a selection protocol, we may view the possibility of rushing as a strategy for players. That is, a player may choose to wait until others have played, and only then submit his message. Thus, NE does not suffice as a solution concept for such games. However, if only one player plays in each round, then this does not matter (since rushing is only allowed *within* rounds), and so for such protocols NE is a reasonable solution concept. For general protocols, we would like the protocol to be optimal for players *regardless of the order of play within a round*. This motivates the following.

For any q -round protocol \mathcal{P} , we can construct protocol \mathcal{P}' with at most qn rounds, and such that only one player sends a message in each round. We say that \mathcal{P}' is a *linearization* of \mathcal{P} , and it is constructed as follows: Let $\pi : [n] \times [q] \mapsto [nq]$ be some bijective map. Then \mathcal{P}' is such that in round ℓ , if $(i, k) = \pi^{-1}(\ell)$ then player i sends a message sampled from S_i^k . This is well-defined for oblivious protocols, and essentially means players play in an arbitrary order, but only one player per round. For non-oblivious protocols, we require π to be *round-respecting*: $\pi(i, k) = \ell$ if and only if for all $j \in [n]$ and $k' < k$ it holds that $\pi(j, k') < \ell$. That is, here the arbitrary ordering is only within rounds.

We note that the idea of considering all linearizations appears also in Ong et al. [29].

Our first solution concept for selection protocols in the full-information model is a full-information dominant strategy equilibrium, which essentially means that for any player i , regardless of the messages sent by others in all rounds, i can never strictly increase his utility by deviating from the protocol. The following generalizes the definition of [5] to multi-round protocols.

Definition 3 (full-information dominant strategy equilibrium). *An oblivious selection protocol \mathcal{P} is in a dominant strategy equilibrium if for all type profiles, all linearizations of \mathcal{P} are in a NE.*

An alternative, more direct but equivalent formulation is the following:

Definition 4 (full-information dominant strategy equilibrium – alternative formulation). *An oblivious n -player, q -round, m -item selection protocol is in a full-information dominant strategy equilibrium if for all $i \in [n]$ and messages $M_{-i} = \{M_k^l : k \in [n] \setminus \{i\}, l \in [q]\}$ sent by all other players in all rounds, it holds that $u_i(t_i, f(M_{-i}, M_i)) = u_i(t_i, f(M_{-i}, M_i^l))$, where $M_i, M_i^l \in \text{supp}(S_i^1) \times \dots \times \text{supp}(S_i^q)$.*

Remark 1. The reason we have equality above, as opposed to an inequality, is that the actions in the support of S_i^j are *all* dominant. That is, all these actions are best-responses, even conditioned on the actions of others. It can thus not be that one such action is better than the other, for then the other would not be dominant.

The definition of a full-information dominant strategy equilibrium is rather strong, but still achievable (for example, Theorems [2], [3], and [4] below). We note that our impossibility result, Theorem [1], applies even to our weaker solution concepts.

In a full-information ex post NE the requirement is a bit relaxed: a player i can not strictly increase his expected utility in any round j by deviating, regardless of the messages of players in all rounds up to *and including* round j . That is, regardless of the order of play within the current round, i has no incentive to deviate (on expectation over play in future rounds). The following definition is new:

Definition 5 (full-information ex post Nash equilibrium). *A selection protocol \mathcal{P} is in an ex post NE if for all type profiles, all round-respecting linearizations of \mathcal{P} are in a NE.*

An alternative, more direct but equivalent formulation for this solution concept is a bit more involved, and appears in the full version of this paper [19].

2.4 Rationality in the Face of an Adversary – Definitions

A protocol that satisfies the definitions of Section 2.3 is an optimal strategy for players assuming all others also follow the protocol. If some of the players are adversarial, however, then this may not hold. In this case, we actually want a stronger guarantee. To this end, we need the following definition, first defined by [1] (for normal-form games):

Definition 6 (v -tolerant NE). *A v -tolerant NE in an extensive-form game is a mixed strategy for every player at every node that he owns, such that the following holds: for any $V \subset [n]$ of size at most v , if all players in $[n] \setminus \{V\}$ play their NE strategy, then none of them can obtain a higher expected payoff by deviating from the NE at any of their nodes regardless of the actions of players in V .*

The ideal faithfulness guarantee that we would like for selection protocols is roughly the following: no player should be able to strictly improve his expected payoff by deviating, assuming most players follow the protocol, some play arbitrarily, and the order within any round is also arbitrary.

Definition 7 (full-information v -tolerant ex post NE). *A leader election protocol \mathcal{P} is in a full-information v -tolerant ex post NE if all round-respecting linearizations of \mathcal{P} are in a v -tolerant NE.*

One possible weakening of this definition is to consider an adversary who does not act arbitrarily, but also has his own utility function u_A . Suppose an adversary corrupts a set V of players. Then we say he is playing a *coalition-optimal strategy* with respect to strategies $S = (S_1, \dots, S_n)$ if, when the players not in V follow strategies S_{-V} , the members of V play a coordinated strategy that maximizes the expectation of u_A . We say he is playing a *strictly coalition-optimal strategy* with respect to strategies S if the above holds, and if, at every node owned by some $i \in V$, i follows S_i if his part of the coordinated deviation does not strictly increase the expectation of u_A . (A more formal definition appears in the full version of this paper [19]).

Definition 8 (*v*-tolerant NE with (strictly) self-interested adversary). A *v*-tolerant NE with self-interested adversary in an extensive-form game is a mixed strategy S_j for every player j for every node that he owns, such that the following holds: for any $V \subset [n]$ of size at most v and any player $i \notin V$, if the players in V play any coalition-optimal strategy and the others play their S_j strategy, then i can not increase his expected utility by deviating from S_i . If this holds only when the players V play a strictly coalition-optimal strategy, then the equilibrium is a *v*-tolerant NE with strictly self-interested adversary.

Definition 9 (full-information *v*-tolerant ex post NE with (strictly) self-interested adversary). A leader election protocol \mathcal{P} is in a full-information *v*-tolerant ex post NE with a (strictly) self-interested adversary if all round-respecting linearizations of \mathcal{P} are in a *v*-tolerant NE with a (strictly) self-interested adversary.

2.5 Resilience to Rational Coalitions – Definitions

In Section 2.4 the adversarial coalition could act arbitrarily, or by maximizing some joint utility function u_A . In this section we define notions of rational coalitions – i.e. coalitions that rational players might reasonably want to form. In the following definitions, we assume there is some prescribed protocol \mathcal{P} for the players. When we say players are “at least as well off” or “strictly gain”, this is with respect to following the prescribed protocol.

Definition 10 (Pareto coalition). A coalition V is a Pareto coalition if there exists a coordinated rushing strategy S_V^* for the players in V such that all players in V are at least as well off when playing S_V^* , and one player strictly gains.

Definition 11 (strong coalition). A coalition V is strong if there exists a coordinated rushing strategy S_V^* for players V such that the expected utility of every $i \in V$ strictly increases when playing S_V^* .

Definition 12 (stable coalition). A coalition V is stable if there exists a coordinated rushing strategy S_V^* for players V such that the expected utility of every $i \in V$ strictly increases when playing S_V^* , and, in addition, for all sub-coalitions $V' \subset V$ and any coordinated rushing strategy $S_{V'}^*$, playing $S_{V'}^*$ does not increase the expected utility of all players in V' when players $V \setminus V'$ play S_V^* .

3 Impossibility with Complex Preferences

A player in a selection protocol has *complex preferences* if for any two outcomes $o \neq o'$ he strictly prefers one over the other. We now show that there are no faithful selection protocols for players with such preferences.

Theorem 1. No selection protocol can be in a full-information ex post NE for players with complex preferences.

Proof. Suppose there exists an *oblivious* selection protocol \mathcal{P} in an ex post Nash equilibrium, and fix some round-respecting linearization of \mathcal{P} . Let T be the corresponding game tree, where some player i owns a node u (that is reached with positive probability) at the lowest non-leaf level ℓ . Suppose the protocol specification is for i to play mixed strategy S_i at level ℓ . Now, if different actions in $\text{supp}(S_i)$ result in leaves with different outcomes, then i prefers one outcome over the others (due to complex preferences). However, due to the full-information ex post NE this can not be the case: a player's different actions should not affect his expected utility, for otherwise he would have a beneficial deviation. We conclude that player i 's actions do not influence the final choice of item. Hence, u can safely be omitted, resulting in a new, smaller tree. We continue shrinking the tree in this manner, yielding a deterministic selection protocol (a contradiction). The extension to non-oblivious protocols appears in the full version of this paper [19].

4 Rational Leader Election Protocols

4.1 Basic Faithful Leader Election Protocols

Because of Theorem 1, we must limit the preferences in order to obtain protocols. One natural setting for leader election is that of *self-interested* players: players care only about whether or not they are elected (they either want to win or want to not win), but are indifferent otherwise. Note that if a leader election protocol is in a full-information dominant strategy equilibrium for self-interested players, then the messages sent by others determine whether a player is elected or not (because the equilibrium holds for all type profiles). That player can only determine who is elected if he is not. The same holds for leader election protocols in a full-information ex post NE, but on expectation over messages in future rounds.

There are some basic protocols that we will use in our constructions. The first is a 1-round leader election protocol that is in a full-information dominant strategy equilibrium (but is not resilient). This protocol was given by Antonakopoulos [6] for the uniform distribution, and then generalized by Altman and Tennenholtz [5].

Theorem 2 ([5]). *For any $n \geq 4$ and any distribution \mathcal{D} over $[n]$ there exists a 1-round, n -player leader election protocol \mathcal{P}_{AT} in a full-information dominant strategy equilibrium, and in which each player i is elected with probability $\mathcal{D}(i)$.*

[5] also showed that there is no faithful 1-round leader election protocol for 3 players. The following protocol, which we will use in our constructions, does work for 3 players, albeit at the cost of having 2 rounds⁵. Fix any natural number $k \geq 3$, and denote $i_+ = (i \bmod (k-1)) + 1$. Then for any positive p_1, \dots, p_k with $p_1 + \dots + p_k = 1$ define

⁵ Note that the case of 2 players is impossible by the lower bound of [32], regardless of the number of rounds.

Protocol \mathcal{P}_k

1. Player k chooses one player $i \neq k$, each with probability $\frac{p_{i+}}{1-p_k}$.
2. For each $j \in \{1, \dots, k-1\}$, if player j is chosen in round 1, he elects player k with probability p_k and player j_+ with probability $1-p_k$ as leader.

Proposition 1. \mathcal{P}_k is a leader election protocol in a full-information ex post NE that elects each player i with probability p_i .

4.2 Combining Rationality and Resilience in Leader Election

Neither of the protocols of Section 4.1 is resilient for any $t > 1$. The following theorem can be combined with resilient leader election protocols to obtain protocols that are both resilient and in full-information dominant strategy equilibria.

Theorem 3. For any $n \geq 4$, $k = \Omega(\sqrt{n})$, and any explicit, oblivious $r(n)$ -round leader election protocol \mathcal{P} there exists an explicit protocol \mathcal{P}' in a dominant strategy equilibrium that has $r(\lceil n/4 \rceil)$ rounds. If \mathcal{P} is resilient to $t(n)$ faults, then \mathcal{P}' is resilient to $t(\lfloor n/4 \rfloor) - k$ faults.

Proof. In the protocol below and the rest of the proof, indices are cyclical. We will prove the theorem for n a multiple of 4. The general case follows similar lines. The players are partitioned into 4 disjoint sets C_1, C_2, C_3, C_4 , where $i \in C_j$ if $\lfloor 4i/n \rfloor = j$. The following is done in **parallel**:

1. Each set C_i runs protocol \mathcal{P} to select a representative R_i .
2. For each i , R_i chooses a random player from C_{i+1} , say L_{i+1} , and outputs a random message b_i to \mathcal{P}_{AT} (i.e. b_i is a random element of B_i , where \mathcal{P}_{AT} takes inputs from $B_1 \times \dots \times B_4$).
3. The winner is L_j , where j is the winner of \mathcal{P}_{AT} with inputs b_1, b_2, b_3, b_4 .

Since the 3 steps are done in parallel, all players choose a random player and a random input in step 2., but the output depends only on the choices of the R_i 's.

Fix some player x , and suppose $x \in C_i$. x is chosen as the leader only if R_{i-1} chooses x . The probability that this occurs does not change regardless of the actions of x . Additionally, for x to win, i must be the winner of \mathcal{P}_4 . However, since \mathcal{P}_{AT} is in an ex post NE, no player in C_i can influence the probability that this occurs. Hence, from x 's perspective, it does not matter who is chosen as L_i .

Now consider some cheating coalition of t players. In order for a member of the coalition to win, at least one member of the coalition must be chosen as L_i for some i . In order for this to occur, either R_{i-1} must be a member of the coalition (and then he can choose a fellow member in C_i), or R_{i-1} is an honest player who chooses a member of the coalition. Suppose there are c_1 faulty players in C_{i-1} and c_2 faulty players in C_i , where $c_1 + c_2 \leq t$. Then the probability that there are more than $c_2 + k$ honest players in C_{i-1} who choose a faulty player is at most a constant $e = \exp(-2k^2/(n/4 - c_1)) < \exp(-16k^2/n) < 1$ by a multiplicative Chernoff bound, and using the fact that $c_1 \leq t(n/4) < n/8$ (since

no leader election protocol can be resilient to more than half the players). Thus, with probability at least $1 - e$, there are at most $c_1 + c_2 + k \leq t(n/4)$ players in C_{i-1} who choose a coalition member in C_i . The maximal probability that one of them wins and becomes R_{i-1} is at most a constant $\varepsilon < 1$ (since we can view the honest players who chose a coalition-member as additional faulty players). The probability that a coalition member becomes L_i for any i is thus at most $1 - (1 - \varepsilon)^4 \cdot (1 - e)^4$, which is some constant < 1 .

In Theorem 3 is that the size of the coalition shrinks by about a factor of 4, and so we can not use it to get a faithful protocol with resilience close to the optimal $n/2$. The following protocol has optimal resilience, is in a full-information dominant strategy equilibrium, and has $\log^*(n) + O(1)$ rounds (same as in the state-of-the-art leader election protocols [31,14]). The proof is in the full version of this paper [19].

Theorem 4. *For every constant $\delta > 0$ and $n \geq 4$ there exists an explicit $(\log^* n + O(1))$ -round leader election protocol resilient against $n(1/2 - \delta)$ faults that is in a full-information dominant strategy equilibrium.*

Extensions and Further Results. Theorem 3 can actually be generalized to obtain any distribution over the players. If we plug a 1-round leader election protocol into Theorem 3 with any distribution, we get a 1-round protocol that implements any distribution and is in a full-information dominant strategy equilibrium. This confirms a conjecture of Altman and Tennenholtz [4] about the existence of such protocols in which all players influence the outcome of the protocol in some instance. We can also construct protocols that satisfy a stronger notion of resilience against adversarial coalitions – namely, they have bounded cheaters’ edge [6] – that are in a full-information ex post NE. Finally, our protocols can also be used to construct leader election protocols in which a player is elected at random *from the set of players who want to be elected*. All these extensions can be found in the full version of this paper [19].

4.3 Rationality in the Face of an Adversary

While the protocols of Section 4.2 are resilient against adversarial behavior, they are in equilibrium only if all players follow the protocol. What if this is not the case? Can an honest player’s protocol specification be optimal even when some others play adversarially? The main difficulty here is that a player’s actions may now also influence the actions of adversarial players in future rounds. Even if the protocol is oblivious, an adversary’s strategy might not be. Definition 7 defines a the concept of an full-information v -tolerant ex post NE to deal precisely with this issue.

Unfortunately, Theorem 5 below implies that no leader election protocol can be in a v -tolerant ex post NE, and so we must look for some relaxation. For Definition 7 we make no assumptions about the adversary. If we assume that the adversary also has some preferences, then we may be able to weaken this restriction. We will assume here that the adversary’s goal is to maximize the

probability that some member of his coalition gets elected (the standard assumption for leader election) – that is, we consider Definition 9, where u_A is the probability that a member of the coalition gets elected. Theorem 5 also shows that this relaxation does not suffice:

Theorem 5. *There does not exist a leader election protocol in a v -tolerant ex post Nash equilibrium with self-interested adversary for any $v \geq 1$.*

Proof. Fix some protocol \mathcal{P} in an ex post Nash equilibrium and a round-respecting linearization of \mathcal{P} . Suppose i is the first player who has a mixed strategy in the game, where two possible messages in i 's support are I_1, I_2 . Because i is eventually chosen by \mathcal{P} with some probability that i can not influence himself (he wins with the same expected probability whether he plays I_1 or I_2), there must exist some other player whose choice of messages does influence this probability. In the subtree rooted at the node following i choosing I_1 there must exist some player j who has a strategy S_1 that increases the probability of i getting elected, and some other strategy S_2 that decreases this probability. Because \mathcal{P} is in an ex post Nash equilibrium, these choices of player j do not harm his own chance of getting elected. A valid (adversarial) strategy for player j is to play S_1 whenever i plays I_1 . Alternatively, j can play S_2 whenever i plays I_1 . Because i does not know which strategy j is using (since j is adversarial), and in either case one of I_1 or I_2 is strictly better than the other, no single strategy of player i can be optimal in both cases.

If we limit the adversary even more by assuming that deviation is costly, we can get an explicit protocol. The following assumes that the adversary is strictly self-interested – he is self-interested, but also only deviates if he strictly gains from doing so (a formal definition appears in the full version of this paper [19]).

Theorem 6. *For any positive k and $n = 3^k$ there exists an explicit n -player $2 \log_3(n)$ -round leader election protocol \mathcal{P} that is in a full-information n -tolerant ex post Nash equilibrium with a strictly self-interested adversary. Furthermore, \mathcal{P} is resilient against $n^{\log_3(2)}/2$ faults.*

To get an idea for the proof, we show that \mathcal{P}_k with $k = 3$ and the uniform distribution is in a full-information 3-tolerant ex post NE with a strictly self-interested adversary. If none or all of the players are adversarial, then all non-adversarial players should follow the protocol (since it is in a full-information ex post NE). If two players are adversarial, then they can always force a win, and so the third player may as well follow the protocol. Finally, if only one player is adversarial, then he can not increase his chance of winning (by the full-information ex post NE), and since the adversary is strictly self-interested he will not deviate. Hence, it is also a full-information ex post NE for the others to follow the protocol.

To generalize this to more players, we divide the players into sets of 3, each running \mathcal{P}_3 . We then repeat this on the winners, until only one is left. The full version of this paper [19] contains further details and an analysis of the resilience of this protocol.

4.4 Resilience to Rational Coalitions

In Sections 4.2 and 4.3 the adversary corrupts some set of v players, and coordinates their actions. Here we let players form a “rational coalition” to benefit all members – namely, we consider the definitions of Section 2.5. For the following theorems (whose proofs appear in the full version [19]), we restrict ourselves to the case in which players are self-interested, and all *want* to be elected. First, we show that it is impossible to have resilience against our weakest notion of a rational coalition.

Theorem 7. *Every leader election protocol in a full-information ex post NE has a Pareto coalition of two players.*

For a stronger notion, however, we can get a protocol that side-steps the impossibility of leader election with adversarial coalitions of size $n/2$:

Theorem 8. *There exists an explicit 2-round leader election protocol in a full-information ex post Nash equilibrium with only 1 stable coalition. The coalition is of size $n - 1$.*

The protocol that achieves this is \mathcal{P}_k with $k = n$ and the uniform distribution (see Section 4.1). We also have the following theorem, as a weak illustration that we gained something by weakening our requirement from strong to stable coalitions.

Theorem 9. *For any n -player leader election protocol in a full-information ex post Nash equilibrium, all coalitions of size $n - 1$ are strong.*

5 Rational Random Sampling Protocols

We consider some universe of m items, and will construct protocols that output each item with probability proportional to the number of players who (claim to) like that item most. In the full version [19] we discuss generalizations to other distributions. Due to Theorem 1, we restrict ourselves to *single-minded* players – each i 's type $t_i \in [m]$ is the item he prefers, and he is indifferent about the others. Theorem 1 also implies that no random sampling is possible with $m = 2$. If $m > 2$ but players prefer only one of two items we are sampling from two items. So we must limit the type profiles. We do this by considering balanced profiles: a profile (t_1, \dots, t_n) for m items is z -balanced if each type occurs between $n/m + z$ and $n/m - z$ times.

Theorem 10. *For any $n \geq 66$ and explicit $r(n)$ -round leader election protocol resilient to $t(n)$ faults in a full-information ex post NE, there exists an explicit $(r(n) + 3)$ -round random selection protocol for a universe of size $m \geq 66$ that is in a full-information ex post NE for all $(n(1/66 - 1/m))$ -balanced type profiles. For such profiles, the random selection protocol is resilient to $t(\lfloor n/3 - n/66 \rfloor)$ faults.*

Proof. The protocol is the following:

1. Each player announces his preferred item. Players are split into 3 categories C_1, C_2, C_3 as follows: all players with the same announced type are in the same category, and the categories are “roughly” balanced: sets of players with the same declared type are greedily assigned to the smallest C_i . Fix $c_i = |C_i|$, and note that $|c_i - c_j| \leq d$ for $d = n/66$ (assuming at most one player lies about his preferred item).
2. For each i , players in C_i run the leader election protocol \mathcal{P} to elect a representative R_i .
3. For each i , the players in $C_{i+1} \cup C_{i+2}$ run the leader election protocol \mathcal{P} , and the winner chooses a uniformly random player L_i from C_i .
4. R_1, R_2 , and R_3 run \mathcal{P}_3 . The protocol is run so that players are elected with probabilities $\frac{c_1}{n}$, $\frac{c_2}{n}$, and $\frac{c_3}{n}$ respectively.
5. The protocol’s output is the announced item of player L_j , where j is the winner of \mathcal{P}_3 in the last round.

If a player $i \in C_j$ truthfully announces his type, then he can no longer change the probability of his type getting chosen: he only affects which of the other types are potential winners (via his choice of L_k for $k \neq j$) or which player from C_i participates in \mathcal{P}_3 . However, since \mathcal{P}_3 is in a full-information ex post NE, this does not matter either.

It remains to show that it is optimal for i to truthfully reveal his type. Suppose i ’s preferred item is B , the fraction of other players who announce B is β , and they all get placed in C_j . Suppose i lies about his type and gets placed in $C_k \neq C_j$. How can i cheat? i wins the leader election protocol of step (2) with probability $1/c_k$ and the leader election protocol of step (3) (choosing L_j) with probability $1/(c_k + c_\ell)$ for $\ell \neq j, k$. If i is elected in both leader election protocols, he can force the winner to be a player who wants B with probability at most 1. If he wins only the leader election protocol of step (2), he can cause j to win in \mathcal{P}_3 with probability $1 - c_k/n$. If he wins only the leader election protocol of step (3), he can force L_j to be a player who wants B (but that player wins \mathcal{P}_3 with probability c_j/n). The probability that B is the chosen type given that i is cheating is

$$\begin{aligned} \Pr[B \text{ wins}] &< \left(1 - \frac{1}{c_k + c_\ell} - \frac{1}{c_k} + \frac{1}{c_k} \cdot \frac{1}{c_k + c_\ell}\right) \cdot \beta \\ &\quad + \frac{1}{c_k + c_\ell} \cdot \frac{c_j}{n} + \frac{1}{c_k} \cdot \frac{\beta n}{c_j} \left(1 - \frac{c_k}{n}\right) + \frac{1}{c_k} \cdot \frac{1}{c_k + c_\ell} \\ &= \beta - \frac{\beta}{c_k + c_\ell} - \frac{\beta}{c_k} + \frac{c_j}{(c_k + c_\ell)n} + \frac{\beta n}{c_k \cdot c_j} - \frac{\beta}{c_j} + \frac{1}{c_k} \cdot \frac{1}{c_k + c_\ell}. \end{aligned}$$

By our balancedness assumption, we know that $n/3 - d \leq c_1, c_2, c_3 \leq n/3 + d$, and that $\beta \leq d/n$. Plugging in these values (and performing some manipulations) yields

$$\Pr[B \text{ wins}] < \beta + \left(\frac{n + 3d}{2n - 6d}\right) \frac{1}{n} + \frac{9d + 18}{\left(\frac{n}{3} - d\right)^2}.$$

It can be verified that when $d \leq n/66$ and $n \geq 66$, we get that $\Pr[B \text{ wins}] < \beta + 1/n$. Now, if player i were to bid truthfully, then the probability that B wins would be $\beta + 1/n$ (since i 's vote adds to B 's chance of winning). Thus, it is an optimal strategy for i to bid truthfully.

What about resilience? Suppose there is an adversary of size at most $t(\lfloor n/3 - n/66 \rfloor)$ faults. In order to force an outcome in some predefined set, the adversary must win at least one of the 6 runs of the leader election protocol \mathcal{P} , and each runs on a set of at least $\lfloor n/3 - n/66 \rfloor$ players. Since \mathcal{P} is resilient for this number of adversaries, the probability that the adversary loses all of them is at least ε^6 for some constant $\varepsilon > 0$.

The following works for smaller m , and is proved in the full version [19].

Theorem 11. *For $n \geq 3$, any explicit $r(n)$ -round leader election protocol resilient up to $t(n)$ faults in a full-information ex post NE, and any constant natural number $m \geq 3$, there exists an explicit $(r(n) + 4)$ -round random selection protocol for a universe of size m that is in a full-information ex post NE for all z -balanced profiles, where $z = n/10m^2$. For such profiles, the random selection protocol is resilient up to $t(\lfloor n/m - z \rfloor)$ faults.*

6 Conclusion and Open Problems

Perhaps the main insight of this paper is that the full-information model is a setting that allows for a relatively clean examination of the interplay between rationality and adversarial behavior in the presence of asynchronous communication. While we have explored numerous aspects of this interplay, we are now faced with many more open questions.

The first set of questions consists of direct extensions of the results presented here. For example, can one generalize the types of preferences for which there are faithful and resilient protocols? For random selection protocols, for example, one might consider a setting in which each player likes some set of items, and dislikes the others. Are there random sampling protocols with weaker balancedness assumptions? How about such protocols that are rational in the face of an adversary, or resilient to rational coalitions? Also, are there protocols with few strong coalitions? Finally, one may consider approximate solution concepts: for example, one may desire all linearizations of a protocol to be in an ε -Nash equilibrium for a small but positive ε . Note that in this case our impossibility result of Theorem 1 no longer applies.

The second set of questions is more open-ended. What can one say about rationality for more general protocol problems in the full-information model? And are there other tractable models for the study of the interplay between rationality and adversarial behavior?

Acknowledgements. I would like to thank Ran Canetti, Moni Naor, and Omer Reingold for helpful conversations. I am also grateful to the anonymous referees for their comments.

References

1. Abraham, I., Dolev, D., Gonen, R., Halpern, J.: Distributed computing meets game theory: Robust mechanisms for rational secret sharing and multiparty computation. In: Proceedings of 25th Annual ACM Symposium on Principles of Distributed Computing, pp. 53–62 (2006)
2. Aiyer, A., Alvisi, L., Clement, A., Dahlin, M., Martin, J.-P., Porth, C.: Bar fault tolerance for cooperative services. In: Proceedings of 20th ACM Symposium on Operating Systems Principles, pp. 45–58 (2005)
3. Alon, N., Naor, M.: Coin-flipping games immune against linear-sized coalitions. *SIAM Journal of Computing* 22(2), 403–417 (1993)
4. Altman, A., Tennenholtz, M.: Selection games and deterministic lotteries (2008), <http://iew3.technion.ac.il/~moshet/selection-lottery.pdf>
5. Altman, A., Tennenholtz, M.: Strategyproof deterministic lotteries under broadcast communication. In: Proceedings of AAMAS (2008)
6. Antonakopoulos, S.: Fast leader-election protocols with bounded cheaters' edge. In: Proceedings of STOC, pp. 187–196 (2006)
7. Aumann, R.J.: Acceptable points in general cooperative n -person games. *Contributions to the Theory of Games, Annals of Mathematical Studies IV*, 287–324 (1959)
8. Ben-Or, M., Linial, N.: Collective coin flipping. *Advances in Computing Research* 5, 91–115 (1989)
9. Bernheim, B.D., Peleg, B., Whinston, M.: Coalition proof nash equilibrium: Concepts. *Journal of Economic Theory* 42(1), 1–12 (1989)
10. Brandt, F., Fischer, F., Shoham, Y.: On strictly competitive multi-player games. In: Proceedings of AAAI (2006)
11. Damgård, I.B.: Interactive hashing can simplify zero-knowledge protocol design without computational assumptions. In: Stinson, D.R. (ed.) CRYPTO 1993. LNCS, vol. 773, pp. 100–109. Springer, Heidelberg (1994)
12. Damgård, I.B., Goldreich, O., Wigderson, A.: Hashing functions can simplify zero-knowledge protocol design (too). Technical Report TR RS-94-39, BRICS (1994)
13. Ding, Y.Z., Harnik, D., Rosen, A., Shaltiel, R.: Constant-round oblivious transfer in the bounded storage model. In: Proceedings of STOC (2004)
14. Feige, U.: Noncryptographic selection protocols. In: Proceedings of 40th Annual Symposium on Foundations of Computer Science, pp. 142–152 (1999)
15. Feigenbaum, J., Shenker, S.: Distributed algorithmic mechanism design: Recent results and future directions. In: Proceedings of the 6th International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications, pp. 1–13 (2002)
16. Goldreich, O., Goldwasser, S., Linial, N.: Fault-tolerant computation in the full information model. *SIAM J. Computing* 27(2) (1998)
17. Goldreich, O., Sahai, A., Vadhan, S.: Honest-verifier statistical zero-knowledge equals general statistical zero-knowledge. In: Proceedings of 30th STOC (1998)
18. Gradwohl, R.: Fault tolerance in distributed mechanism design. In: Proceedings of the 4th Workshop on Internet and Network Economics (2008)
19. Gradwohl, R.: Rationality in the Full-Information Model, <http://www.kellogg.northwestern.edu/faculty/Gradwohl/html/papers/rle.pdf>
20. Gradwohl, R., Vadhan, S., Zuckerman, D.: Random selection with an adversarial majority. In: Dwork, C. (ed.) CRYPTO 2006. LNCS, vol. 4117, pp. 409–426. Springer, Heidelberg (2006)

21. Halpern, J.: A computer scientist looks at game theory. *Games and Economic Behavior* 45(1), 114–131 (2003)
22. Halpern, J.: Computer science and game theory: A brief survey. *The New Palgrave Dictionary of Economics* (2008)
23. Kalai, E.: Large robust games. *Econometrica* 72(6), 1631–1665 (2004)
24. Katz, J.: Bridging game theory and cryptography: Recent results and future directions. In: Canetti, R. (ed.) *TCC 2008*. LNCS, vol. 4948, pp. 251–272. Springer, Heidelberg (2008)
25. Monderer, D., Tennenholtz, M.: Distributed games: from mechanisms to protocols. In: *Proceedings of the 16th National Conference on Artificial Intelligence*, pp. 32–37 (1999)
26. Moreno, D., Wooders, J.: Coalition-proof equilibrium. *Games and Economic Behavior* 17(1), 80–112 (1996)
27. Naor, M., Ostrovsky, R., Venkatesan, R., Yung, M.: Perfect zero-knowledge arguments for NP can be based on general complexity assumptions. *J. Cryptology* 11 (1998)
28. Ong, S.J., Parkes, D., Rosen, A., Vadhan, S.: Fairness with an honest minority and a rational majority. Preliminary version (October 2007)
29. Ong, S.J., Parkes, D., Rosen, A., Vadhan, S.: Fairness with an honest minority and a rational majority. In: Reingold, O. (ed.) *TCC 2009*. LNCS, vol. 5444, pp. 36–53. Springer, Heidelberg (2009)
30. Osborne, M.J., Rubinstein, A.: *A Course in Game Theory*. MIT Press, Cambridge (1994)
31. Russell, A., Zuckerman, D.: Perfect-information leader election in $\log^* n + O(1)$ rounds. *Journal of Computer and System Sciences* 63, 612–626 (2001)
32. Saks, M.: A robust noncryptographic protocol for collective coin flipping. *SIAM Journal on Discrete Mathematics* 2(2), 240–244 (1989)
33. Sanghvi, S., Vadhan, S.: The round complexity of two-party random selection. *SIAM Journal of Computing* 32(2), 523–550 (2008)
34. Shneidman, J., Parkes, D.C.: Specification faithfulness in networks with rational nodes. In: *Proceedings of PODC* (2004)