

Byzantine Agreement with a Rational Adversary

Adam Groce, Jonathan Katz*, Aishwarya Thiruvengadam, and Vassilis Zikas**

Department of Computer Science, University of Maryland
{agroce, jkatz, aish, vzikas}@cs.umd.edu

Abstract. Traditionally, cryptographers assume a “worst-case” adversary who can act arbitrarily. More recently, they have begun to consider *rational* adversaries who can be expected to act in a utility-maximizing way. Here we apply this model for the first time to the problem of *Byzantine agreement* (BA) and the closely related problem of *broadcast*, for natural classes of utilities. Surprisingly, we show that many known results (e.g., equivalence of these problems, or the impossibility of tolerating $t \geq n/2$ corruptions) do not hold in the rational model. We study the feasibility of information-theoretic (both perfect and statistical) BA assuming complete or partial knowledge of the adversary’s preferences. We show that perfectly secure BA is possible for $t < n$ corruptions given complete knowledge of the adversary’s preferences, and characterize when statistical security is possible with only partial knowledge. Our protocols have the added advantage of being more efficient than BA protocols secure in the traditional adversarial model.

1 Introduction

The problem of *Byzantine agreement* (BA) was introduced by Lamport, Shostak, and Pease [15] as an abstraction of their earlier work on distributed computation among fallible processors [18]. The problem comes in two flavors, called *consensus* and *broadcast*. In consensus, we have n players each having an input and it is required that they all agree on an output y (consistency), where if all correct players have the same input x then $y = x$ (correctness). In broadcast, only one player (the *sender*) has input, and the requirements are that all players should agree on an output y (consistency), such that if the sender correctly follows the protocol then y is equal to its input (correctness).

In the original work of Lamport et al. [15] BA was motivated by the so-called *Byzantine generals problem*: The generals of the Byzantine army, along with their troops, have encircled an enemy city. Each general is far away from the rest and messengers are used for communication. The generals must agree upon a common plan (to attack or to retreat), though one or more of the generals may be traitors who will attempt to foil the plan. The good generals do not know who the traitors are. If the good generals agree upon the plan unanimously, the

* Supported in part by NSF grants #5-23126 and #5-24541.

** Supported in part by a fellowship from the Swiss National Science Foundation (Project No. PBEZP2-134445).

plan will succeed. The traitors, however, may choose to coordinate in a manner that would mislead the good generals into disagreement.

In the formal definition of the problem, the destructive behavior of the traitors is modeled by assuming that they are corrupted by a (central) adversary who coordinates them and tries to break the security. Breaking the security of BA corresponds to violating some of the aforementioned properties, i.e, correctness or consistency (or both). But imagine that the generals had some additional preferences. Say, for example, the traitors wanted to cause inconsistency among the honest generals but, if they can't do that, they would at least prefer the honest generals retreat rather than attack. Or maybe they just want as few generals to attack as possible. In short, it is realistic to assume that the traitors are not acting arbitrarily, but instead have a clear set of preferences, and prefer some outcomes over others.

In this paper, we model the above by assuming *rational* adversaries. A rational adversary prefers a particular outcome of the protocol and may deviate in an attempt to achieve its preference. This is different from the traditional malicious adversary setting, wherein the only goal of the adversary is to break the security of a protocol. We investigate feasibility of rational Byzantine agreement (RBA) under various assumptions regarding the adversary's preferences. Interestingly, in addition to providing a conceptually simple way of capturing realistic situations like the one described above, the model yields significant differences with respect the traditional (non-rational) BA setting. In particular, many properties — even some well-established impossibility results — that are taken for granted in the traditional model are no longer true in the rational setting.

1.1 Our Results

We present, for the first time, a definition of Byzantine agreement taking into account rational behavior on the part of the adversary. In our work, we adopt a somewhat different approach than that taken in some other work blending game theory and cryptography (see below): rather than treating *all* players as rational, we assume that some players are honest and will follow the protocol without question, while other players (those controlled by the adversary) are rational and will attempt to alter the outcome so as to increase their utility.

We study rational broadcast and Byzantine agreement for a natural class of adversarial utility functions defined by the adversary's preferences over the possible outcomes: agreement on 0, agreement on 1, and disagreement. Interestingly, many of the statements that are considered self-evident in the BA literature break down in the rational setting. Examples include the impossibility of consensus for $t \geq n/2$, the usefulness of setups for statistical (and computational) security, as well as the reduction of consensus to broadcast for $t < n/2$. We also study of feasibility of RBA for all possible orderings on the adversary's preferences in the following two cases: (1) the utility function of the adversary is known, and (2) only the adversary's preference between agreement and disagreement is known (but among the possible outcomes for agreement, it is not known which one is more preferred).

1.2 Related Work

Byzantine agreement and broadcast have been studied extensively, and we limit ourselves to a discussion of the main results. Early work showed that (without any setup), Byzantine agreement is possible if and only if the number of corrupted parties t is strictly less than $1/3$ of the total number of parties n . The situation changes when a trusted setup allowing digital signatures (e.g., a public-key infrastructure (PKI)) is assumed. Such a setup does not make a difference for perfect security as there is always some (possibly negligible) probability that the adversary guesses the secret keys of the honest parties and breaks the security of the BA protocol. However, as shown in [18], computationally secure broadcast can be achieved for arbitrary $t < n$ if it is assumed that honest players can sign the messages they send. The same bound was shown to be achievable for the case of statistical security, assuming a setup for information-theoretically secure (pseudo-)signatures [4, 19]. It follows from the definition (for a traditional, worst-case adversary) that consensus is impossible in any setting when $t \geq n/2$.

There has recently been a significant amount of interest in bridging cryptographic and game-theoretic models and definitions; see, e.g., [12, 13, 10, 1, 17, 3, 11]. We refer to [14] for a (now slightly outdated) survey. Most closely related to our own work, the BAR model [2] was developed to capture Byzantine, altruistic, and rational behavior in a distributed system; protocols that tolerate a combination of Byzantine and rational nodes were proposed for reliable broadcast [6], state machine replication, and gossip-based multicast [16]. In contrast to these works, our model considers some nodes to be rational and the rest to be honest. In work done concurrently with our own, Bei et al. [5] have shown that rational consensus is impossible in the presence of colluding rational agents and crash failures. Their model assumes that each player is individually rational, and wants to strategically manipulate the protocol according to his own set of preferences. In addition, some of these agents could have crash failures. This is incomparable to our model where we assume that some players honestly follow the protocol while the rest are under the control of a centralized, rational adversary.

2 Byzantine Agreement

We briefly review the traditional definitions of broadcast and consensus. We let P_1, \dots, P_n denote the parties running the protocol, and let t be (an upper bound on) the number of deviating parties. We let v_i and w_i denote the input and output, respectively, of P_i . We assume the standard network model, where all pairs of players have (authenticated) point-to-point channels. We assume synchronous communication and allow a computationally unbounded adversary. We refer to a (static) adversary who corrupts up to t parties as a t -adversary.

Definition 1 (Consensus) *Each player P_i initially has input v_i . A protocol is a perfectly secure consensus protocol if it satisfies the following properties:*

1. (Consistency): *All honest players output the same value w .*

2. (Correctness): If all honest players begin with the same input value, i.e. $v_i = v$ for all i , then $w = v$.

Definition 2 (Broadcast) We refer to player P_1 as the sender, who is transmitting his input v_1 to the remaining $n - 1$ receivers. A protocol is a perfectly secure broadcast protocol if it satisfies the following properties:

1. (Consistency): All honest players output the same value w .
2. (Correctness): If the sender P_1 is honest then $w = v_1$.

The definitions of BA for statistical and computational security are obtained by requiring that the corresponding properties are satisfied except with negligible probability in the presence of an unbounded and a computationally bounded adversary, respectively.

Detectable Broadcast. A useful primitive in our constructions is *detectable broadcast*, which was defined by Fitzi et al. in [7] as a relaxation of the definition of broadcast. Informally, a detectable broadcast guarantees that at the protocol termination, either a successful realization of broadcast has been achieved or all honest parties have agreed that the protocol has been aborted.

The formal definition for detectable broadcast [7] is as follows. For simplicity, we describe the definition for the case where the input is a bit.

Definition 3 (Detectable broadcast) A protocol for detectable broadcast must satisfy the following properties:

1. (Correctness): All honest players either abort or accept and output 0 or 1. If any honest player aborts, so does every honest player. If no honest players abort, then the output satisfies the security conditions of broadcast (according to Definition 2).
2. (Completeness): If all players are honest, all players accept (and therefore achieve broadcast without error).
3. (Fairness): If any honest player aborts then the adversary receives no information about the sender's bit.

A protocol for detectable broadcast was presented in [8] that satisfies the above definition except with some negligible error probability in the presence of an unbounded adversary (i.e., with statistical security) corrupting arbitrary many parties ($t < n$). The protocol for detectable broadcast given by [8] requires $t + 5$ rounds and $O(n^8(\log n + k)^3)$ total bits of communication, where k is a security parameter and $t < n$.

3 Rational Byzantine Agreement

We next define our model of a rational adversary and within it the definitions of rational BA. A *rational adversary* is characterized by some utility function which describes his preference over possible outcomes of the protocol execution. In the

following we describe generic definitions of security in the presence of such an adversary; subsequently, we specify a natural class of utilities for an adversary attacking a BA protocol. Towards the end of the section, we also study the relation between the traditional and the rational definition of BA.

The adversary’s utility In any analysis of security against rational adversaries, one needs to define the adversaries’ behavior. The first step to doing this is to define their utility, which provides a method for deciding which outcomes an adversary (or any other rational player) prefers to which others. We present a definition of utility that we believe is natural, reasonable, and can be worked with easily. In particular, we consider real utility, i.e., the utility is described by real numbers associated with particular outcomes. For simplicity, we limit ourselves to protocols that are attempting to broadcast or agree on a single bit. The adversary’s utility is defined on the following events: (1) All honest players output (agree on) 1, (2) all honest players output (agree on) 0, and (3) honest players have disagreeing output. In particular we define the utility function of the adversary as follows: For values $u_0, u_1, u_2 \in \mathbb{R}$:

$$U[\text{agreement on } 0] := u_0, \quad U[\text{agreement on } 1] := u_1, \quad \text{and} \quad U[\text{disagreement}] := u_2$$

For simplicity we assume that the values u_0, u_1 , and u_2 are distinct, but all our proofs go through even if some of them are equal. We assume that rational players will choose from the strategies available to them the one that results in the most preferred outcome. However, since strategies and the protocol can be randomized, a particular set of strategies will imply not a particular outcome but a particular distribution over outcomes. The utility of a distribution is then the expected value of the utility of an outcome drawn from that distribution.

Definition 4 (Utility) *An expectation utility is a utility that conforms to the following condition. Using D_z to represent the probability distribution where outcome $z \in Z$ occurs with probability 1, we require that $U(D) = E[U(D_z)|z \leftarrow D]$.*

The above utility function corresponds, of course, to a substantial simplification of possible outcomes. For example, some sorts of disagreement could be preferred over any unanimous output while other types are disliked. Nevertheless, these outcomes capture a meaningful portion of potential outcome variation. In order to maximize the strength of our results, we assume that the adversary knows the inputs of the honest players (which are disclosed very early in most protocols anyway) and can therefore choose its strategy to maximize utility for that particular input set.

Definition. We assume that all corrupted players are colluding. Equivalently, there is a single adversary that directs the actions of up to t (non-adaptively) corrupted players. The other players are honest, meaning that rather than acting according to their selfish interests they simply run the protocol as specified. This means that the “game” we are considering actually only has one player. We are essentially considering what is a Nash equilibrium strategy for the adversary.

However, the Nash equilibrium of a one-player game is simply a straightforward utility-optimization, so we leave out the complexities of Nash equilibria in our definition. When we refer to a “strategy” we mean simply a function that takes as input the view of the adversary so far and outputs its next message/action.

Definition 5 (Perfect security) *A protocol for broadcast or consensus is perfectly secure against rational adversaries controlling t players with utility U if for every t -adversary there is a strategy S such that for any choice of input for honest players*

1. (*S is tolerable*): S induces a distribution of final outputs D in which no security condition is violated with nonzero probability, and
2. (*S is Nash*): For any strategy $S' \neq S$ with induced output distribution D' : $U(D) \geq U(D')$.

In addition to this standard notion, we will be considering a definition following from statistical equilibria. Here we introduce a security parameter k . The strategy sees the security parameter at the beginning of the game and can alter its behavior based on that parameter. We require not that the security-respecting strategy be perfectly optimal but that it is within a negligible distance to optimal. This means that the incentive to deviate could be made arbitrarily small, and would get extremely small very quickly as the security parameter is raised.

Definition 6 (Statistical security) *A protocol for broadcast or consensus is statistically secure against rational adversaries controlling t players with utility U if for every t -adversary there is a strategy S such that for any choice of input for honest players S induces a distribution of final outputs D_k when the security parameter is k and the following properties hold:*

1. (*S is tolerable*): no security condition is violated with nonzero probability in D_k for any k , and
2. (*S is statistical Nash*): for any strategy $S' \neq S$ with induced output distributions D'_k there is a negligible function $\text{negl}(\cdot)$ such that $U(D_k) + \text{negl}(k) > U(D'_k)$.

Remark (Statistical tolerability and honestly perfect protocols). The above definition requires that the strategy S is *perfectly* tolerable. One could weaken this definition to require *statistical tolerability*, i.e., require that the tolerability property is satisfied except with some negligible probability. We argue that this does not make a difference for any protocol which, assuming no party is corrupted, satisfies the properties of BA with perfect security (we refer to such protocols as *honestly perfect*). Indeed, for an honestly perfect protocol there exists a strategy S^H , namely the strategy corresponding to honestly executing the protocol, which is perfectly tolerable. Let D_k^H denote the distribution induced by S^H and D_k denote the distribution induced by the optimal strategy S from the above definition where we require that S is only statistically tolerable. The statistical tolerability of S implies that $U(D_k) = U(D_k^H) \pm \text{negl}(k)$. This, combined with the fact that S is statistical Nash, implies that $U(D_k^H) + \text{negl}(k) \geq U(D'_k)$ for all D'_k . Hence, D_k^H is statistically Nash and perfectly tolerable.

We note that a *computational* security definition could also be considered. Such a definition is equivalent to the statistical case except that the strategy function is required to be computable in polynomial time (in k). We do not consider computational security in this work. However, all our statements about feasibility with statistical security hold also for computational security.

Relation to the Traditional Definition. In the following we show that rational BA reduces to traditional BA. The proof is based on the observation that if a protocol is secure according to the traditional definition of BA, then in RBA every adversarial strategy is Nash.

Theorem 1 (BA implies RBA) *If protocol Π perfectly securely realizes traditional consensus (resp., broadcast) in the presence of a (non-rational) t -adversary, then Π is perfectly secure for consensus (resp., broadcast) against rational t -adversaries with utility U . The statement holds also for statistical security assuming the protocol π is honestly perfect.¹*

4 Rational Byzantine Agreement: Basic Results

In this section, we motivate the study of feasibility of rational BA by demonstrating that some of the results that are taken for granted in the traditional BA literature become invalid in the rational setting.

The Traditional Impossibility of Consensus Fails. It is well-known that when $t \geq n/2$, there exists no consensus protocol which tolerates a t -adversary, even when the parties have access to a broadcast channel. The idea of the proof is the following: Consider the setting where the first $n/2$ of the parties have input 0, and the remaining have input 1. Assume the following adversarial scenarios: (A) the adversary corrupts the first $n/2$ or (B) the adversary corrupts the last $n/2$ parties; in both scenarios the adversary has the corrupted parties execute their correct protocol. In Scenario A, the honest parties should all output 1, whereas in Scenario B they should output 0. Consider now a third scenario (Scenario C) where the adversary does not corrupt any party. Because this Scenario is indistinguishable from Scenario B, the first half of parties should output 0; however, because Scenario C is indistinguishable from Scenario A, the second half of parties should output 1, which leads to contradiction.

We show that in the rational setting this impossibility does not, in general, hold: Consider a rational adversary with utility $u_2 > u_1 > u_0$. Then, as the following lemma suggests, assuming a (traditional) broadcast channel, there exists a consensus protocol tolerating arbitrary many parties, i.e., $t < n$, even with perfect security. The protocol, denoted as Π' works as follows:

¹ Recall that a BA protocol is honestly perfect if it satisfies the perfect security definition in the absence of an adversary.

Protocol $\Pi'(v_1, \dots, v_n)$

1. Every party P_i broadcasts his input v_i .
2. If all parties broadcast the same value then output it, otherwise output 0.

The idea of the proof is that the adversary will never try to introduce an inconsistency, as if he does so he will be punished with his worst preferred outcome (i.e., 0).

Lemma 1 *The protocol Π' described above is (perfectly) secure for consensus against rational t -adversaries with $t < n$ and utilities satisfying $u_2 > u_1 > u_0$.*

The Traditional Reduction of Consensus to Broadcast Fails. Traditional consensus and broadcast are known to be equivalent assuming $t < n/2$ parties are corrupted. The idea is the following: assuming consensus, broadcast can be achieved by the sender sending his input to every party and then invoking consensus on the received values. Similarly, assuming broadcast (and $t < n/2$) consensus can be achieved by having every party broadcast his input and taking the majority of the broadcasted values he receives to be his output.

Surprisingly, the above straightforward reduction of consensus to broadcast does not transfer through to the rational setting. Informally, the reason for the failure of the reduction is the inherent incomposability issue that appears in most rational security models. In particular, for the case of the above reduction, it is possible that when attacking a consensus protocol that uses broadcast protocols as subroutines, that an adversary can achieve a *desired* outcome in the consensus protocol by violating the security of the broadcast subroutines in ways that would seem, on their own, *undesirable*.

Due to space limitations, we refer the reader to the full version of this work for a concrete description of our counterexample and the corresponding analysis.

Luckily, the reduction in the other direction is still successful. The proof can be found in the full version of this paper.

Theorem 2 *Assume that a consensus protocol exists that is secure against rational adversaries with a particular utility. A protocol can be constructed for broadcast that is secure against rational adversaries with the same utility.*

Equivalence of Statistical and Perfect Security (Setup-Independent) Perhaps one of the most unexpected differences between traditional and rational BA is the fact that in the rational setting (with real utilities), a setup does not offer anything with respect to feasibility in the information theoretic setting, as perfect security is possible for $t < n$. This is in contrast to the traditional BA where a setup is known to bring the exact bound for statistical complexity from $t < n/3$ (for both consensus and broadcast) down to $t < n/2$ for consensus and $t < n$ for broadcast. The following theorem states that the two levels of information theoretic security, i.e., perfect and statistical security, are equivalent in the rational setting. The idea for reducing perfect to statistical security is the following: because the values u_0 , u_1 , and u_2 are real numbers, in any statistical protocol one can fix the security parameter to be large enough, so that the adversary does

not any more have an incentive to cheat, which will lead to a perfectly secure protocol. A proof can be found in the full version of this work.

Theorem 3 *There exists a statistically secure protocol for rational consensus (resp., broadcast) tolerating some adversary \mathcal{A} if and only if there exists a perfectly secure protocol for rational consensus (resp., broadcast) tolerating \mathcal{A} .*

5 Feasibility Assuming Complete Knowledge

In this section, we give a complete characterization of feasibility of RBA for information-theoretic security. Note that as implied by Theorem 3, the bound for statistical and perfect security is the same. In fact, this bound is $t < n$ independent of the adversary's preference. This is stated in the following theorem. Due to space limitations, we only describe the idea of the proof and sketch the main argument for some of the cases, and refer to the full version of this work for a complete handling of all the cases.

Theorem 4 *There exists a protocol for perfectly secure Byzantine agreement tolerating a rational t -adversary, where the utilities $u_0, u_1, u_2 \in \mathbb{R}$ are known, tolerating arbitrarily many corruptions, i.e., $t < n$. The statement holds both for broadcast and consensus.*

Proof (sketch). The general idea is, as in the proof of Lemma 1, to force the adversary play the strategy which guarantees the security of the protocol by having the protocol punish him in case he does not. Note that, by Theorem 2, it suffices to describe consensus protocols; furthermore, by Theorem 3, it suffices to achieve statistical security. The proof considers two cases: (1) the adversary's most favorable choice is *not* disagreement, and (2) the adversary's most favorable choice is disagreement. In the first case, the following consensus protocol works:

1. Every player P_i sends his input v_i to every player P_j .
2. For every P_j : if the same value was received from every P_i then output it, otherwise output the bit b' which the adversary prefers least (i.e., $u_{b'} < \min\{u_2, u_{1-b'}\}$).

Intuitively, the above protocol is secure, as when there is pre-agreement among the honest parties, then the adversary has no incentive to destroy it, and when there is disagreement they will output b' .

The somewhat more involved setting occurs when the top choice of the adversary is to force disagreement. We consider the following consensus protocol:

1. Each P_i uses detectable broadcast [7] to broadcast his input v_i .
2. If any abort occurs or there is disagreement among the broadcasted value, then output the adversary's least preferred bit b' (i.e., $u_{b'} < \min\{u_2, u_{1-b'}\}$). Otherwise output the value broadcasted by all parties.

The fact that the above protocol is (statistically) secure is argued as follows: Consistency follows trivially from the consistency of detectable broadcast; furthermore, because the adversary has no incentive to break the detectable broadcast protocol, and, by the security of detectable broadcast, when it does not abort it satisfies the correctness property, the following adversarial strategy is a Nash equilibrium: allow all honest senders to broadcast their input and have the corrupted senders broadcast $1 - b'$.

6 Feasibility with Partial Knowledge

So far we have been considering only cases where the protocol is designed with full knowledge of the adversary's preferences, but it is also possible to consider cases where the adversary's preferences are not fully known. The goal is to guarantee security against any adversary that has preferences consistent with the limited information that we have. If no information about the adversary's preferences is known, this reduces to the traditional setting of a malicious adversary. If some information exists, however, the situation can be more interesting.

In the full-information settings we have been considering up to this point, statistical and perfect security are provably equivalent (Theorem 3). This result does not hold in the case of partial information. In fact, we can give protocols and impossibility results that prove that no such equivalence holds when setup is allowed. Similarly, we give results that show that when no setup is allowed, consensus and broadcast are not equivalent. In order to show these results, we consider the situation where it is known whether or not the adversary wishes to create disagreement between the parties, but it is not known what the adversary's preferences are among different potential agreeing outputs.

Disagreement is the Adversary's Most Favorable Option. If we consider the setting where disagreement is known to be the adversary's most preferred outcome, all the impossibility proofs from the traditional world apply. We can therefore deduce that the bounds for both broadcast and consensus are the same as in the traditional setting. We state this formally below and refer the reader to the full version of this work for the proof. (This is a tight bound, since it matches the possibility result from the malicious setting, which of course also applies in any rational setting.)

Theorem 5 *Assuming $n \geq 3$, there does not exist a perfectly or statistically secure rational consensus protocol that tolerates any t -adversary with $t \geq n/3$ and disagreement as the most-preferred outcome.*

Disagreement is the Adversary's Least Favorable Option. Finally, we consider the case where the adversary wants to avoid disagreement, but has unknown preferences on the other outcomes. This case is interesting because it provides an instance where what is possible is provably different for broadcast than it is for consensus, in contrast to the traditional setting of a malicious adversary

(with perfect security). The proof of the following theorem can be found in the full version.

Theorem 6 *Assuming $n \geq 3$, there exists a perfectly secure rational consensus protocol tolerating any t -adversary with disagreement as the least-preferred outcome if and only if $t < n/2$. The statement holds also for statistical security.*

We complete our characterization by looking at feasibility of broadcast for the case where disagreement is the adversary's least-preferred outcome. As shown in the following theorem, in that case broadcast can be achieved, by the trivial multi-send protocol, tolerating an arbitrary number of corruptions. Combined with Theorem 6, this proves that in this setting perfectly secure broadcast is easier to achieve than consensus.

Theorem 7 *There exists a perfectly secure rational broadcast protocol tolerating any t -adversary for $t < n$ with disagreement as the least-preferred outcome. The statement holds for statistical security as well.*

Proof. We use the following perfectly secure protocol: The sender sends his input to every party who outputs the value received from the sender. If $t = n - 1$ the security conditions are trivially satisfied. For the case where there are at least two honest players, we consider two cases. In the first case, the sender is honest. As a result, all honest players are sent the correct output and no error is made. In the second case, the sender is not honest. In this case, the adversary would not have the sender send disagreeing messages to honest parties, since disagreement is the least-preferred outcome. However, because the sender is dishonest, any agreeing output from the honest parties is consistent with the security conditions, so no security violation can occur. Statistical security follows immediately, since it is a weaker definition.

Acknowledgments

We thank Dov Gordon for collaboration during the early stages of this work [9].

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 *PODC 2006*, pages 53–62. ACM Press, 2006.
2. A. S. Aiyer, L. Alvisi, A. Clement, M. Dahlin, J.-P. Martin, and C. Porth. BAR fault tolerance for cooperative services. In *SOSP 2005*, pages 45–58. ACM, 2005.
3. G. Asharov, R. Canetti, and C. Hazay. Towards a game theoretic view of secure computation. In *Advances in Cryptology — Eurocrypt 2011*, volume 6632 of *LNCS*, pages 426–445. Springer, 2011.
4. B. Baum-Waidner, B. Pfitzmann, and M. Waidner. Unconditional Byzantine agreement with good majority. In *STACS '91*, volume 480 of *LNCS*, pages 285–295. Springer, 1991.

5. X. Bei, W. Chen, and J. Zhang. Distributed consensus resilient to both crash failures and strategic manipulations, 2012. arXiv 1203.4324.
6. A. Clement, H. C. Li, J. Napper, J.-P. Martin, L. Alvisi, and M. Dahlin. BAR primer. In *DSN 2008*, pages 287–296. IEEE Computer Society, 2008.
7. M. Fitzi, N. Gisin, U. Maurer, and O. von Rotz. Unconditional Byzantine agreement and multi-party computation secure against dishonest minorities from scratch. In *Advances in Cryptology — Eurocrypt 2002*, volume 2332 of *LNCS*, pages 482–501. Springer, 2002.
8. M. Fitzi, D. Gottesman, M. Hirt, T. Holenstein, and A. Smith. Detectable Byzantine agreement secure against faulty majorities. In *PODC 2002*, pages 118–126. ACM Press, 2002.
9. S. D. Gordon and J. Katz. Byzantine agreement with a rational adversary. Rump session presentation, Crypto 2006.
10. S. D. Gordon and J. Katz. Rational secret sharing, revisited. In *SCN 2006*, volume 4116 of *LNCS*, pages 229–241. Springer, 2006.
11. A. Groce and J. Katz. Fair computation with rational players. In *Advances in Cryptology — Eurocrypt 2012*, volume 7237 of *LNCS*, pages 81–98. Springer, 2012.
12. J. Halpern and V. Teague. Rational secret sharing and multiparty computation: Extended abstract. In *STOC 2004*, pages 623–632. ACM Press, 2004.
13. S. Izmalkov, S. Micali, and M. Lepinski. Rational secure computation and ideal mechanism design. In *FOCS 2005*, pages 585–595. IEEE Computer Society Press, 2005.
14. J. Katz. Bridging game theory and cryptography: Recent results and future directions (invited talk). In *TCC 2008*, volume 4948 of *LNCS*, pages 251–272. Springer, 2008.
15. L. Lamport, R. E. Shostak, and M. C. Pease. The Byzantine generals problem. *ACM Trans. Programming Language Systems*, 4(3):382–401, 1982.
16. H. C. Li, A. Clement, E. L. Wong, J. Napper, I. Roy, L. Alvisi, and M. Dahlin. Bar gossip. In *OSDI 2006*, pages 191–204. USENIX Association, 2006.
17. S. Ong, D. Parkes, A. Rosen, and S. Vadhan. Fairness with an honest minority and a rational majority. In *TCC 2009*, volume 5444 of *LNCS*, pages 36–53. Springer, 2009.
18. M. Pease, R. E. Shostak, and L. Lamport. Reaching agreement in the presence of faults. *Journal of the ACM*, 27(2):228–234, 1980.
19. B. Pfitzmann and M. Waidner. Unconditional Byzantine agreement for any number of faulty processors. In *STACS '92*, volume 577 of *LNCS*, pages 339–350. Springer, 1992.