Lower Bounds on Implementing Robust and Resilient Mediators

Ittai Abraham[1], Danny Dolev[2,⋆], and Joseph Y. Halpern[3,⋆⋆]

[1] Hebrew University
ittaia@cs.huji.ac.il
[2] Hebrew University
dolev@cs.huji.ac.il
[3] Cornell University
halpern@cs.cornell.edu

Abstract. We provide new and tight lower bounds on the ability of players to implement equilibria using cheap talk, that is, just allowing communication among the players. One of our main results is that, in general, it is impossible to implement three-player Nash equilibria in a bounded number of rounds. We also give the first rigorous connection between Byzantine agreement lower bounds and lower bounds on implementation. To this end we consider a number of variants of Byzantine agreement and introduce reduction arguments. We also give lower bounds on the running time of two player implementations. All our results extended to lower bounds on (k, t)-*robust* equilibria, a solution concept that tolerates deviations by coalitions of size up to k and deviations by up to t players with unknown utilities (who may be malicious).

1 Introduction

The question of whether a problem in a multiagent system that can be solved with a trusted mediator can be solved by just the agents in the system, without the mediator, has attracted a great deal of attention in both computer science (particularly in the cryptography community) and game theory. In cryptography, the focus on the problem has been on *secure multiparty computation*. Here it is assumed that each agent i has some private information x_i. Fix functions $f_1, \ldots, f_n$. The goal is have agent i learn $f_i(x_1, \ldots, x_n)$ without learning anything about x_j for $j \neq i$ beyond what is revealed by the value of $f_i(x_1, \ldots, x_n)$. With a trusted mediator, this is trivial: each agent i just gives the mediator its private value x_i; the mediator then sends each agent i the value $f_i(x_1, \ldots, x_n)$. Work on

[⋆] Part of the work was done while the author visited Cornell university. The work was funded in part by ISF, ISOC, NSF, CCR, and AFOSR.

[⋆⋆] Supported in part by NSF under grants CCR-0208535, ITR-0325453, and IIS-0534064, by ONR under grant N00014-01-10-511, by the DoD Multidisciplinary University Research Initiative (MURI) program administered by the ONR under grants N00014-01-1-0795 and N00014-04-1-0725, and by AFOSR under grant FA9550-05-1-0055.

R. Canetti (Ed.): TCC 2008, LNCS 4948, pp. 302–319, 2008.

multiparty computation (see [18] for a survey) provides conditions under which this can be done. In game theory, the focus has been on whether an equilibrium in a game with a mediator can be implemented using what is called *cheap talk*—that is, just by players communicating among themselves (see [28] for a survey).

There is a great deal of overlap between the problems studied in computer science and game theory. But there are some significant differences. Perhaps the most significant difference is that, in the computer science literature, the interest has been in doing multiparty computation in the presence of possibly malicious adversaries, who do everything they can to subvert the computation. On the other hand, in the game theory literature, the assumption is that players have preference and seek to maximize their utility; thus, they will subvert the computation iff it is in their best interests to do so. Following [1], we consider here both rational adversaries, who try to maximize their utility, and possibly malicious adversaries (who can also be considered rational adversaries whose utilities we do not understand).

1.1 Our Results

In this paper we provide new and optimal lower bounds on the ability to implement mediators with cheap talk. Recall that a *Nash equilibrium* σ is a tuple of strategies such that given that all other players play their corresponding part of σ then the best response is also to play σ. Given a Nash equilibrium σ we say that a strategy profile ρ is a *k-punishment strategy for σ* if, when all but k players play their component of ρ, then no matter what the remaining k players do, their payoff is strictly less than what it is with σ. We now describe some highlights of our results in the two simplest settings: (1) where rational players cannot form coalitions and there are no malicious players (this gives us the solution concept of Nash equilibrium) and (2) where there is at most one malicious player. We describe our results in a more general setting in Section 1.2.

No bounded implementations: In [1] it was shown that any Nash equilibrium with a mediator for three-player games with a 1-punishment strategy can be implemented using cheap talk. The expected running time of the implementation is constant. It is natural to ask if implementations with a bounded number of rounds exist for all three-player games. Theorem 2 shows this is not the case, implementations must have infinite executions and cannot be bounded for all three-player games. This lower bound highlights the importance of using randomization. An earlier attempt to provide a three-player cheap talk implementation [8] uses a bounded implementation, and hence cannot work in general. The key insight of the lower bound is that when the implementation is bounded, then at some point the punishment strategy must become ineffective. The details turn out to be quite subtle. The only other lower bound that we are aware of that has the same flavor is the celebrated FLP result [15] for reaching agreement in asynchronous systems, which also shows that no bounded implementation exists. However, we use quite different proof techniques than FLP.

Byzantine Agreement and Game Theory: We give the first rigorous connection between Byzantine agreement lower bounds and lower bounds on implementation. To get the lower bounds, we need to consider a number of variants of Byzantine agreement, some novel. The novel variants require new impossibility results. We have four results of this flavor:

1. Barany [6] gives an example to show that, in general, to implement an equilibrium with a mediator in a three-player game, it is necessary to have a 1-punishment strategy. Using the power of randomized Byzantine agreement lower bounds we strengthen his result and show in Theorem 4 that we cannot even get an ϵ-implementation in this setting.
2. Using the techniques of [7] or [17], it is easy to show that any four-player game Nash equilibrium with a mediator can be implemented using cheap talk even if no 1-punishment strategy exists. Moreover, these implementations are *universal*; they do not depend on the players' utilities. In Theorem 3 we prove that universal implementations do not exist in general for three-player games. Our proof uses a nontrivial reduction to the weak Byzantine agreement (WBA) problem [24]. To obtain our lower bound, we need to prove a new impossibility result for WBA, namely, that no protocol with a finite expected running time can solve WBA.
3. In [1] we show that for six-player games with a 2-punishment strategy, any Nash equilibrium can be implemented even in the presence of at most one malicious player. In Theorem 5 we show that for five players even ϵ–-implementation is impossible. The proof uses a variant of Byzantine agreement; this is related to the problem of *broadcast with extended consistency* introduced by Fitzi et al. [16]. Our reduction maps the rational player to a Byzantine process that is afraid of being detected and the malicious player to a standard Byzantine process.
4. In Theorem 8, we show that for four-player games with at most one malicious player, to implement the mediator, we must have a PKI setup in place, even if the players are all computationally bounded and even if we are willing to settle for ϵ–implementations. Our lower bound is based on a reduction to a novel relaxation of the Byzantine agreement problem.

Bounds on running time: We provide bounds on the number of rounds needed to implement two-player games. In Theorem 9(a) we prove that the expected running time of any implementation of a two-player mediator equilibrium must depend on the utilities of the game, even if there is a 1-punishment strategy. This is in contrast to the three-player case, where the expected running time is constant. In Theorem 9(b) we prove that the expected running time of any ϵ–implementation of a two-player mediator equilibrium for which there is no 1-punishment strategy must depend on ϵ. Both results are obtained using a new two-player variant of the secret-sharing game. The only result that we are aware of that has a similar spirit is that of Boneh and Naor [9], where it is shown that two-party protocols with "bounded unfairness" of ϵ must have running time that depends on the value of ϵ. The implementations given by Urbano and Vila [31,32] in the two-player case are independent of the utilities; the above results show that their implementation cannot be correct in general.

1.2 Our Results for Implementing Robust and Resistant Mediators

In [1] (ADGH from now on), we argued that it is important to consider deviations by both rational players, who have preferences and try to maximize them, and players that can be viewed as malicious, although it is perhaps better to think of them as rational players whose utilities are not known by the other players or mechanism designer. We considered equilibria that are (k, t)-*robust*; roughly speaking, this means that the equilibrium tolerates deviations by up to k rational players, whose utilities are presumed known, and up to t players with unknown utilities (i.e., possibly malicious players). We showed how (k, t)-robust equilibria with mediators could be implemented using cheap talk, by first showing that, under appropriate assumptions, we could implement secret sharing in a (k, t)-robust way using cheap talk. These assumptions involve standard considerations in the game theory and distributed systems literature, specifically, (a) the relationship between k, t and n, the total number of players in the system; (b) whether players know the exact utilities of other players; (c) whether there are broadcast channels or just point-to-point channels; (d) whether cryptography is available; and (e) whether the game has a $(k + t)$-*punishment strategy*; that is, a strategy that, if used by all but at most $k + t$ players, guarantees that every player gets a worse outcome than they do with the equilibrium strategy. Here we provide a complete picture of when implementation is possible, providing lower bounds that match the known upper bounds (or improvements of them that we have obtained). The following is a high-level picture of the results. (The results discussed in Section 1.1 are special cases of the results stated below. Note that all the upper bounds mentioned here are either in ADGH, slight improvements of results in ADGH, or are known in the literature; see Section 3 for the details. The new results claimed in the current submission are the matching lower bounds.)

- If $n > 3k + 3t$, then mediators can be implemented using cheap talk; no punishment strategy is required, no knowledge of other agents' utilities is required, and the cheap-talk strategy has bounded running time that does not depend on the utilities (Theorem 1(a) in Section 3).
- If $n \leq 3k+3t$, then we cannot, in general, implement a mediator using cheap talk without knowledge of other agents' utilities (Theorem 3). Moreover, even if other agents' utilities are known, we cannot, in general, implement a mediator without having a punishment strategy (Theorem 4) nor with bounded running time (Theorem 2).
- If $n > 2k + 3t$, then mediators can be implemented using cheap talk if there is a punishment strategy (and utilities are known) in finite expected running time that does not depend on the utilities (Theorem 1(b) in Section 3).
- If $n \leq 2k + 3t$, then we cannot, in general, ϵ-implement a mediator using cheap talk, even if there is a punishment strategy and utilities are known (Theorem 5).

- If $n > 2k + 2t$ and we can simulate broadcast then, for all ϵ, we can ϵ-implement a mediator using cheap talk, with bounded expected running time that does not depend on the utilities in the game or on ϵ (Theorem 1(c) in Section 3). (Intuitively, an ϵ-implementation is an implementation where a player can gain at most ϵ by deviating.)
- If $n \leq 2k+2t$, we cannot, in general, ϵ-implement a mediator using cheap talk even if we have broadcast channels (Theorem 7). Moreover, even if we assume cryptography and broadcast channels, we cannot, in general, ϵ-implement a mediator using cheap talk with expected running time that does not depend on ϵ (Theorem 9(b)); even if there is a punishment strategy, then we still cannot, in general, ϵ-implement a mediator using cheap talk with expected running time independent of the utilities in the game (Theorem 9(a)).
- If $n > k + 3t$ then, assuming cryptography, we can ϵ-implement a mediator using cheap talk; moreover, if there is a punishment strategy, the expected running time does not depend on ϵ (Theorem 1(e) in Section 3).
- If $n \leq k + 3t$, then even assuming cryptography, we cannot, in general, ϵ-implement a mediator using cheap talk (Theorem 8).
- If $n > k + t$, then assuming cryptography and that a PKI (Public Key Infrastructure) is in place, [1] we can ϵ-implement a mediator (Theorem 1(d) in Section 3); moreover, if there is a punishment strategy, the expected running time does not depend on ϵ (Theorem 1(e) in Section 3).

The lower bounds are existential results; they show that if certain conditions do not hold, then there exists an equilibrium that can be implemented by a mediator that cannot be implemented using cheap talk. There are other games where these conditions do not hold but we can nevertheless implement a mediator.

1.3 Related Work

There has been a great deal of work on implementing mediators, both in computer science and game theory. The results above generalize a number of results that appear in the literature. We briefly discuss the most relevant work on implementing mediators here. Other work related to this paper is discussed where it is relevant.

In game theory, the study of implementing mediators using cheap talk goes back to Crawford and Sobel [11]. Barany [6] shows that if $n \geq 4$, $k = 1$, and $t = 0$ (i.e., the setting for Nash equilibrium), a mediator can be implemented in a game where players do not have private information. Forges [17] provides what she calls a *universal mechanism* for implementing mediators; essentially, when combining her results with those of Barany, we get the special case of Theorem 1(a) where $k = 1$ and $t = 0$. Ben-Porath [8] considers implementing a mediator with cheap talk in the case that $k = 1$ if $n \geq 3$ and there is a 1-punishment strategy. He seems to have been the first to consider punishment strategies (although his notion is different from ours: he requires that there be an equilibrium that

[1] We can replace the assumption of a PKI here and elsewhere by the assumption that there is a trusted preprocessing phase where players may broadcast.

is dominated by the equilibrium that we are trying to implement). Heller [22] extends Ben-Porath's result to allow arbitrary k. Theorem 1(b) generalizes Ben-Porath and Heller's results. Although Theorem 1(b) shows that the statement of Ben-Porath's result is correct, Ben-Porath's implementation takes a bounded number of rounds; Theorem 2 shows it cannot be correct. [2] Heller proves a matching lower bound; Theorem 5 generalizes Heller's lower bound to the case that $t > 0$. (This turns out to require a much more complicated game than that considered by Heller.) Urbano and Vila [31,32] use cryptography to deal with the case that $n = 2$ and $k = 1$; [3] Theorem 1(e)) generalizes their result to arbitrary k and t. However, just as with Ben-Porath, Urbano and Vila's implementation takes a bounded number of rounds; As we said in Section 1.1, Theorem 9(a) shows that it cannot be correct.

In the cryptography community, results on implementing mediators go back to 1982 (although this terminology was not used), in the context of *(secure) multiparty computation*. Since there are no utilities in this problem, the focus has been on essentially what we call here *t-immunity*: no group of t players can prevent the remaining players from learning the function value, nor can they learn the other players' private values. Results of Yao [33] can be viewed as showing that if $n = 2$ and appropriate computational hardness assumptions are made, then, for all ϵ, we can obtain 1-immunity with probability greater than $1 - \epsilon$ if appropriate computational hardness assumptions hold. Goldreich, Micali, and Wigderson [19] extend Yao's result to the case that $t > 0$ and $n > t$. Ben-Or, Goldwasser, and Wigderson [7] and Chaum, Crépeau, and Damgard [10] show that, without computational hardness assumptions, we can get t-immunity if $n > 3t$; moreover, the protocol of Ben-Or, Goldwasser, and Wigderson does not need an ϵ "error" term. Although they did not consider utilities, their protocol actually gives a (k, t)-robust implementation of a mediator using cheap talk if $n > 3k + 3t$; that is, they essentially prove Theorem 1(a). (Thus, although these results predate those of Barany and Forges, they are actually stronger.) Rabin and Ben-Or [29] provide a t-immune implementation of a mediator with "error" ϵ if broadcast can be simulated. Again, when we add utilities, their protocol actually gives an ϵ–(k, t)-robust implementation. Thus, they essentially prove Theorem 1(c). Dodis, Halevi, and Rabin [12] seem to have been the first to apply cryptographic techniques to game-theoretic solution concepts; they consider the case that $n = 2$ and $k = 1$ and there is no private information (in which case the equilibrium in the mediator game is a *correlated equilibrium* [5]); their result is essentially that of Urbano and Vila [32] (although their protocol does not suffer form the problems of that of Urbano and Vila).

Halpern and Teague [21] were perhaps the first to consider the general problem of multiparty computation with rational players. In this setting, they essentially prove Theorem 1(d) for the case that $t = 0$ and $n \geq 3$. However, their focus is on

[2] Although Heller's implementation does not take a bounded number of rounds, it suffers from problems similar to those of Ben-Porath.

[3] However, they make somewhat vague and nonstandard assumptions about the cryptographic tools they use.

the solution concept of *iterated deletion*. They show that there is no Nash equilibrium for rational multiparty computation with rational agents that survives iterated deletion and give a protocol with finite expected running time that does survive iterated deletion. If $n \leq 3(k + t)$, it follows easily from Theorem 2: that there is no multiparty computation protocol that is a Nash equilibrium, we do not have to require that the protocol survive iterated deletion to get the result if $n \leq 3(k + t)$. Various generalizations of the Halpern and Teague results have been proved. We have already mentioned the work of ADGH. Lysanskaya and Triandopoulos [27] independently proved the special case of Theorem 1(c) where $k = 1$ and $t + 1 < n/2$ (they also consider survival of iterated deletion); Gordon and Katz [20] independently proved a special case of Theorem 1(d) where $k = 1$, $t = 0$, and $n \geq 2$.

In this paper we are interested in implementing equilibrium by using standard communication channels. An alternate option is to consider the possibility of simulating equilibrium by using much stronger primitives. Izmalkov, Micali, and Lepinski [23] show that, if there is a punishment strategy and we have available strong primitives that they call *envelopes* and *ballot boxes*, we can implement arbitrary mediators perfectly (without an ϵ error) in the case that $k = 1$, in the sense that every equilibrium of the game with the mediator corresponds to an equilibrium of the cheap-talk game, and vice versa. In [26,25], these primitives are also used to obtain implementation that is perfectly collusion proof in the model where, in the game with the mediator, coalitions cannot communicate. (By way of contrast, we allow coalitions to communicate.) Unfortunately, envelopes and ballot boxes cannot be implemented under standard computational and systems assumptions [25].

The rest of this paper is organized as follows. In Section 2, we review the relevant definitions. In Section 3, we briefly discuss the upper bounds, and compare them to the results of ADGH. In Section 4, we prove the lower bounds.

2 Definitions

We give a brief description of the definitions needed for our results here. More detailed definitions and further discussion can be found in [3].

We are interested in implementing mediators. Formally, this means we need to consider three games: an *underlying game* Γ, an extension Γ_d of Γ with a mediator, and a cheap-talk extension Γ_{CT} of Γ. Our underlying games are *(normal-form) Bayesian games*. These are games of incomplete information, where players make only one move, and these moves are made simultaneously. The "incomplete information" is captured by assuming that nature makes the first move and chooses for each player i a *type* in some set $\mathcal{T}_i$, according to some distribution that is commonly known. Formally, a Bayesian game Γ is defined by a tuple $(N, \mathcal{T}, A, u, \mu)$, where N is the set of players, $\mathcal{T} = \times_{i \in N} \mathcal{T}_i$ is the set of possible types, μ is the distribution on types, $A = \times_{i \in N} A_i$ is the set of action profiles, and $u_i : \mathcal{T} \times A$ is the utility of player i as a function of the types prescribed by nature and the actions taken by all players.

Given an underlying Bayesian game Γ as above, a game Γ_d with a mediator d that extends Γ is, informally, a game where players can communicate with the mediator and then perform an action from Γ. The utility of player i in Γ_d depends just on its type and the actions performed by all the players. Although we think of a *cheap-talk* game as a game where players can communicate with each other (using point-to-point communication and possibly broadcast), formally, it is a game with a special kind of mediator that basically forwards all the messages it receives to their intended recipients. We assume that mediators and players are just interacting Turing machines with access to an unbiased coin (which thus allows them to choose uniformly at random from a finite set of any size). Γ_{CT} denotes the cheap-talk extension of Γ.

When considering a deviation by a coalition K, one may want to allow the players in K to communicate with each other. If Γ' is an extension of an underlying game Γ (including Γ itself) and $K \subseteq N$, let $\Gamma' + CT(K)$ be the extension of Γ where the mediator provides private cheap-talk channels for the players in K in addition to whatever communication there is in Γ'. Note that $\Gamma_{\mathrm{CT}} + CT(K)$ is just Γ_{CT}; players in K can already talk to each other in Γ_{CT}.

A *strategy* for player i in a Bayesian game Γ is a function from i's type to an action in A_i; in a game with a mediator, a strategy is a function from i's type and message history to an action. We allow behavior strategies (i.e., randomized strategies); such a strategy gets an extra argument, which is a sequence of coin flips (intuitively, what a player does can depend on its type, the messages it has sent and received if we are considering games with mediators, and the outcome of some coin flips). We use lower-case Greek letters such as σ, τ, and ρ to denote a strategy profile; σ_i denotes the strategy of player i in strategy profile σ; if $K \subseteq N$, then σ_K denotes the strategies of the players in K and σ_{-K} denotes the strategies of the players not in K. Given a strategy profile σ a player $i \in N$ and a type $t_i \in T_i$ let $u_i(t_i, \sigma)$ be the expected utility of player i given that his type is t_i and each player $j \in N$ is playing the strategy σ_j. Note that a strategy profile—whether it is in the underlying game, or in a game with a mediator extending the underlying game (including a cheap-talk game)—induces a mapping from type profiles to distributions over action profiles. If Γ_1 and Γ_2 are extension of some underlying game Γ, then strategy σ_1 in Γ_1 *implements* a strategy σ_2 in Γ_2 if both σ and σ' induce the same function from types to distributions over actions. Note that although our informal discussion in the introduction talked about *implementing mediators*, the formal definitions (and our theorems) talk about implementing strategies. Our upper bounds show that, under appropriate assumptions, for *every* (k, t)-robust equilibrium σ in a game Γ_1 with a mediator, there exists an equilibrium σ' in the cheap-talk game Γ_2 corresponding to Γ_1 that implements σ; the lower bounds in this paper show that, if these conditions are not met, there exists a game with a mediator and an equilibrium in that game that cannot be implemented in the cheap-talk game. Since our definition of games with a mediator also allow arbitrary communication among the agents, it can also be shown that every equilibrium in a cheap-talk game can be implemented in the mediator game: the players simply ignore the mediator and communicate with each other.

The utility function in the games we consider is defined on type and action profiles. Note that we use the same utility function both for an underlying game Γ and all extensions of it. As usual, we want to talk about the expected utility of a strategy profile, or of a strategy profile conditional on a type profile. We abuse notation and continue to use u_i for this, writing for example, $u_i(t_K, \sigma)$ to denote the expected utility to player i if the strategy profile σ is used, conditional on the players in K having the types t_K. Since the strategy σ here can come from the underlying game or some extension of it, the function u_i is rather badly overloaded. We sometimes include the relevant game as an argument to u_i to emphasize which game the strategy profile σ is taken from, writing, for example, $u_i(t_K, \Gamma', \sigma)$.

We now define the main solution concept used in this paper: (k, t)-robust equilibrium. The k indicates the size of coalition we are willing to tolerate, and the t indicates the number of players with unknown utilities. These t players are analogues of faulty players or adversaries in the distributed computing literature, but we can think of them as being perfectly rational. Since we do not know what actions these t players will perform, nor do we know their identities, we are interested in strategies for which the payoffs of the remaining players are immune to what the t players do.

Definition 1. *A strategy profile σ in a game Γ is t-*immune *if, for all $T \subseteq N$ with $|T| \leq t$, all strategy profiles τ, all $i \notin T$, and all types $t_i \in \mathcal{T}_i$ that occur with positive probability, we have $u_i(t_i, \Gamma + CT(T), \sigma_{-T}, \tau_T) \geq u_i(t_i, \Gamma, \sigma)$.*

Intuitively, σ is t-immune if there is nothing that players in a set T of size at most t can do to give the remaining players a worse payoff, even if the players in T can communicate.

Our notion of (k, t)-robustness requires both t-immunity and the fact that, no matter what t players do, no subset of size at most k can all do better by deviating, even with the help of the t players, and even if all $k + t$ players share their type information.

Definition 2. *Given $\epsilon \geq 0$, σ is an ϵ–(k, t)-*robust equilibrium *in game Γ if σ is t-immune and, for all $K, T \subseteq N$ such that $|K| \leq k$, $|T| \leq t$, and $K \cap T = \emptyset$, and all types $t_{K \cup T} \in \mathcal{T}_{K \cup T}$ that occur with positive probability, it is not the case that there exists a strategy profile τ such that*

$$u_i(t_{K \cup T}, \Gamma + CT(K \cup T), \tau_{K \cup T}, \sigma_{-(K \cup T)}) > u_i(t_i, \Gamma + CT(T), \tau_T, \sigma_{-T}) + \epsilon$$

for all $i \in K$. A (k, t)-robust equilibrium is just a 0–(k, t)-robust equilibrium.

Note that a $(1, 0)$-robust equilibrium is just a Nash equilibrium, and an ϵ–$(1, 0)$-robust equilibrium is what has been called an ϵ-Nash equilibrium in the literature. The notion $(k, 0)$-robust equilibrium is essentially Aumann's [4] notion of resilience to coalitions, except that we allow communication by coalition members (see [3] for a discussion of the need for such communication). Heller [22] used essentially this notion. The notion $(0, t)$-robustness is somewhat in the spirit of Eliaz's [13] notion of t fault-tolerant implementation. Both our notion of $(0, t)$-robustness and Eliaz's notion of t-fault tolerance require that what the players not in T do is a best response to whatever the players in T do (given that all the players not in T follow the recommended strategy); however, Eliaz does not require an analogue of t-immunity.

In [1] we considered a stronger version of robust equilibrium. Roughly speaking, in this stronger version, we require that, if a coalition deviates, only one coalition member need be better off, rather than all coalition members. In [3] we formally define this stronger notion and discuss its motivation. We note that all our lower and upper bounds works for both notions; we focus on Definition 1 here because it is more standard in the game theory literature. (Other notions of equilibrium have been considered in the literature; see the appendix for discussion.)

In this paper, we are interested in the question of when a (k, t)-robust equilibrium σ in a game Γ_d with a mediator extending an underlying game Γ can be implemented by an ϵ–(k, t)-robust equilibrium σ' in the cheap-talk extension Γ_{CT} of Γ. If this is the case, we say that σ' is an *ϵ–(k, t)-robust* implementation of σ. (We sometimes say that $(\Gamma_{\mathrm{CT}}, \sigma')$ is an *ϵ–(k, t)-robust* implementation of (Γ_d, σ) if we wish to emphasize the games.)

3 The Possibility Results

Definition 3. *If Γ_d is an extension of an underlying game Γ with a mediator d, a strategy profile ρ in Γ is a* k-punishment strategy with respect to a strategy profile σ in Γ_d *if for all subsets $K \subseteq N$ with $|K| \leq k$, all strategies ϕ in $\Gamma + CT(K)$, all types $t_K \in T_K$, and all players $i \in K$:*

$$u_i(t_K, \Gamma_d, \sigma) > u_i(t_K, \Gamma + CT(K), \phi_K, \rho_{-K}).$$

If the inequality holds with $\geq$ replacing $>$, ρ is a weak *k-punishment strategy with respect to σ.*

Intuitively, ρ is k-punishment strategy with respect to σ if, for any coalition K of at most k players, even if the players in K share their type information, as long as all players not in K use the punishment strategy in the underlying game, there is nothing that the players in K can do in the underlying game that will give them a better expected payoff than playing σ in Γ_d.

The notion of utility variant is used to make precise that certain results do not depend on knowing the players' utilities (see [3] for details).

Theorem 1. *Suppose that Γ is Bayesian game with n players and utilities u, d is a mediator that can be described by a circuit of depth c, and σ is a (k, t)-robust equilibrium of a game Γ_d with a mediator d.*

(a) If $3(k + t) < n$, then there exists a strategy σ_{CT} in $\Gamma_{\mathrm{CT}}(u)$ such that for all utility variants $\Gamma(u')$, if σ is a (k, t)-robust equilibrium of $\Gamma_d(u')$, then $(\Gamma_{\mathrm{CT}}(u'), \sigma_{\mathrm{CT}})$ implements $(\Gamma_d(u'), \sigma)$. The running time of σ_{CT} is $O(c)$.

(b) If $2k + 3t < n$ and there exists a $(k + t)$-punishment strategy with respect to σ, then there exists a strategy σ_{CT} in Γ_{CT} such that σ_{CT} implements σ. The expected running time of σ_{CT} is $O(c)$.

(c) If $2(k + t) < n$ and broadcast channels can be simulated, then, for all $\epsilon > 0$, there exists a strategy $\sigma_{\mathrm{CT}}^{\epsilon}$ in Γ_{CT} such that $\sigma_{\mathrm{CT}}^{\epsilon}$ ϵ-implements σ. The running time of $\sigma_{\mathrm{CT}}^{\epsilon}$ is $O(c)$.

(d) If $k + t < n$ then, assuming cryptography and that a PKI is in place, there exists a strategy $\sigma_{\mathrm{CT}}^{\epsilon}$ in Γ_{CT} such that $\sigma_{\mathrm{CT}}^{\epsilon}$ ϵ-implements σ. The expected running time of $\sigma_{\mathrm{CT}}^{\epsilon}$ is $O(c) \cdot f(u) \cdot O(1/\epsilon)$ where $f(u)$ is a function of the utilities.

(e) If $k + 3t < n$ or if $k + t < n$ and a trusted PKI is in place, and there exists a $(k + t)$-punishment strategy with respect to σ, then, assuming cryptography, there exists a strategy $\sigma_{\mathrm{CT}}^{\epsilon}$ in Γ_{CT} such that $\sigma_{\mathrm{CT}}^{\epsilon}$ ϵ-implementers σ. The expected running time of $\sigma_{\mathrm{CT}}^{\epsilon}$ is $O(c) \cdot f(u)$ where $f(u)$ is a function of the utilities but is independent of ϵ.

We briefly comment on the differences between Theorem 1 and the corresponding Theorem 4 of ADGH. In ADGH, we were interested in finding strategies that were not only (k, t)-robust, but also survived iterated deletion of weakly dominated strategies. For part (a), in ADGH, a behavioral strategy was used that had no upper bound on running time. This was done in order to obtain a strategy that survived iterated deletion. However, it is observed in ADGH that, without this concern, a strategy with a known upper bound can be used. As we observed in the introduction, part (a), as stated, actually follows from [7]. Part (b) here is the same as in ADGH. In part (c), we assume here the ability to simulate broadcast; ADGH assumes cryptography. As we have observed, in the presence of cryptography, we can simulate broadcast, so the assumption here is weaker. In any case, as observed in the introduction, part (c) follows from known results [29]. Parts (d) and (e) are new, and will be proved in [2]. The proof uses ideas from [19] on multiparty computation. For part (d), where there is no punishment strategy, ideas from [14] on getting *ϵ-fair* protocols are also required. Our proof of part (e) shows that if $n > k + 3t$, then we can essentially set up a PKI on the fly. These results strengthen Theorem 4(d) in ADGH, where punishment was required and n was required to be greater than $k + 2t$.

4 The Impossibility Results

No Bounded Implementations

We prove that it is impossible to get an implementation with bounded running time in general if $2k + 3t < n \leq 3k + 3t$. This is true even if there is a punishment strategy. This result is optimal. If $3k + 3t < n$, then there does exist a bounded implementation; if $2k + 3t < n \leq 3k + 3t$ there exists an unbounded implementation that has constant *expected* running time.

Theorem 2. *If $2k + 3t < n \leq 3k + 3t$, there is a game Γ and a strong (k, t)-robust equilibrium σ of a game Γ_d with a mediator d that extends Γ such that there exists a $(k+t)$-punishment strategy with respect to σ for which there do not exist a natural number c and a strategy σ_{CT} in the cheap talk game extending Γ such that the running time of σ_{CT} on the equilibrium path is at most c and σ_{CT} is a (k, t)-robust implementation of σ.*

Proof. We first assume that $n = 3$, $k = 1$, and $t = 0$. We consider a family of 3-player games $\Gamma_3^{n,k+t}$, where $2k+3t < n \leq 3k+3t$, defined as follows. Partition $\{1, \ldots, n\}$ into three sets B_1, B_2, and B_3, such that B_1 consists of the first $\lfloor n/3 \rfloor$ elements in $\{1, \ldots, n\}$, B_3 consists of the last $\lceil n/3 \rceil$ elements, and B_2 consists of the remaining elements.

Let p be a prime such that $p > n$. Nature chooses a polynomial f of degree $k + t$ over the p-element field $GF(p)$ uniformly at random. For $i \in \{1, 2, 3\}$, player i's type consists of the set of pairs $\{(h, f(h)) \mid h \in B_i\}$. Each player wants to learn $f(0)$ (the secret), but would prefer that other players do not learn the secret. Formally, each player must play either 0 or 1. The utilities are defined as follows:

- if all players output $f(0)$ then all players get 1;
- if player i does not output $f(0)$ then he gets -3;
- otherwise players i gets 2.

Consider the mediator game where each player is supposed to tell the mediator his type. The mediator records all the pairs (h, v_h) it receives. If at least $n-t$ pairs are received and there exists a unique degree $k + t$ polynomial that agrees with at least $n - t$ of the pairs then the mediator interpolates this unique polynomial f' and sends $f'(0)$ to each player; otherwise, the mediator sends 0 to each player.

Let σ_i be the strategy where player i truthfully tells the mediator his type and follows the mediator's recommendation. It is easy to see that σ is a $(1, 0)$-robust equilibrium (i.e., a Nash equilibrium). If a player i deviates by misrepresenting or not telling the mediator up to t of his shares, then everyone still learns; if the player misrepresents or does not tell the mediator about more of his shares, then the mediator sends the default value 0. In this case i is worse off. For if 0 is indeed the secret, which it is with probability $1/2$, i gets 1 if he plays 0, and -3 if he plays 1. On the other hand, if 1 is the secret, then i gets 2 if he plays 1 and -3 otherwise. Thus, no matter what i does, his expected utility is at most $-1/2$. This argument also shows that if ρ_i is the strategy where i decides 0 no matter what, then ρ is a 1-punishment strategy with respect to σ.

Suppose, by way of contradiction, that there is a cheap-talk strategy σ' in the game Γ_{CT} that implements σ such that any execution of σ' takes at most c rounds. We say that a player i *learns the secret by round b of σ'* if, for all executions (i.e., plays) r and r' of σ' such that i has the same type and the same message history up to round b, the secret is the same in r and r'. Since we have assumed that all plays of σ' terminate in at most c rounds, it must be the case that all players learn the secret by round c of σ'. For if not, there are two executions r and r' of σ' that i cannot distinguish by round c, where the secret is different in r and r'. Since i must play the same move in r and r', in one case he is not playing the secret, contradicting the assumption that σ' implements σ. Thus, there must exist a round $b \leq c$ such that all three players learn the secret at round b of σ' and, with nonzero probability, some player, which we can assume without loss of generality is player 1, does not learn the secret at round $b - 1$ of σ'. This means that there exists a type t_1 and message history h_1 for player 1 of length $b - 1$ that occurs with positive probability when player 1 has type t_1 such that, after $b-1$ rounds, if player 1 has type t_1 and history h_1, player 1 considers it possible that the secret could be either 0 or 1. Thus, there must exist type profiles t and t' that correspond to polynomials f and f' such that $t_1 = t'_1$, $f(0) \neq f'(0)$ and, with positive probability, player 1 can have history h_1 with both t and t', given that all three players play σ'.

Let h_2 be a history for player 2 of length $b-1$ compatible with t and h_1 (i.e., when the players play σ', with positive probability, player 1 has h_1, player 2 has h_2, and the true type profile is t); similarly, let h_3 be a history of length $b-1$ for player 3 compatible with t' and h_1. Note that player i's action according to σ_i is completely determined by his type, his message history, and the outcome of his coin tosses. Let $\sigma'_2[t_2, h_2]$ be the strategy for player 2 according to which player 2 uses σ'_2 for the first $b - 1$ rounds, and then from round b on, player 2 does what it would have done according to σ'_2 if its type had been t_2 and its message history for the first $b-1$ rounds had been h_2 (that is, player 2 modifies his actual message history by replacing the prefix of length $b-1$ by h_2, and leaving the rest of the message history unchanged). We can similarly define $\sigma'_3[t'_3, h_3]$. Consider the strategy profile $(\sigma'_1, \sigma'_2[t_2, h_2], \sigma'_3[t'_3, h_3])$. Since $\sigma'_i[t_i, h_i]$ is identical to σ'_i for the first $b - 1$ steps, for $i = 2, 3$, there is a positive probability that player 1 will have history h_1 and type t_1 when this strategy profile is played. It should be clear that, conditional on this happening, the probability that player 1 plays 0 or 1 is independent of the actual types and histories of players 2 and 3. This is because players 2 and 3's messages from time b depend only on i's messages, and not on their actual type and history. Thus, for at least one of 0 and 1, it must be the case that the probability that player 1 plays this value is strictly less than 1. Suppose without loss of generality that the probability of playing $f(0)$ is less than 1.

We now claim that $\sigma'_3[t'_3, h_3]$ is a profitable deviation for player 3. Notice that player 3 receives the same messages for the first b rounds of σ' and $(\sigma'_1, \sigma'_2, \sigma'_3[t'_3, h_3])$. Thus, player 3 correctly plays the secret no matter what the type profile is, and gets payoff of at least 1. Moreover, if the type profile is t, then, by construction, with positive probability, after $b-1$ steps, player 1's history will be h_1 and player 2's history will be h_2. In this case, σ'_2 is identical to $\sigma'_2[t_2, h_2]$, so the play will be identical to $(\sigma'_1, \sigma'_2[t_2, h_2], \sigma'_3[t'_3, h_3])$. Thus, with positive probability, player 1 will not output $f(0)$, and player 3 will get payoff 2. This means player 3's expected utility is greater than 1.

For the general case, suppose that $2k+3t < n \leq 3k+3t$. Consider the n-player game $\Gamma^{n,k,t}$, defined as follows. Partition the players into three groups, B_0, B_1, and B_2, as above. As in the 3-player game, nature chooses a polynomial f of degree $k+t$ over the field $GF(p)$ with a prime $p > n$ uniformly at random, but now player i's type is just the pair $(i, f(i))$. Again, the players want to learn $f(0)$, but would prefer that other players do not learn the secret, and must output a value in F. The payoffs are similar in spirit to the 3-player game:

- if at least $n-t$ players output $f(0)$ then all players that output $f(0)$ get 1;
- if player i does not output $f(0)$ then he gets -3;
- otherwise player i gets 2.

The mediator's strategy is essentially identical to that in the 3-player game (even though now it is getting one pair (h, v_h) from each player rather than a set of such pairs from a single player). Similarly, each player i's strategy in $\Gamma_d^{n,k,t}$, which we denote σ_i^n, is essentially identical to the strategy in the 3-player game with the mediator. Again, if ρ_i^n is the strategy in the n-player game where i plays 0 no matter what his type, then it is easy to check that ρ^n is a $(k+t)$-punishment strategy with respect to σ^n.

Now suppose, by way of contradiction, that there exists a strategy σ' in the cheap-talk extension $\Gamma_{\mathrm{CT}}^{n,k,t}$ of $\Gamma^{n,k,t}$ that is a (k,t)-robust implementation of σ^n such that all executions of σ' take at most c rounds. We show in [3] that we can use σ' to get a $(1,0)$-robust implementation in the 3-player mediator game $\Gamma_{3,d}^{n,k+t}$, contradicting the argument above. □

Byzantine Agreement and Game Theory

In [1] it is shown that if $n > 3k+3t$, we can implement a mediator in a way that does not depend on utilities and does not need a punishment strategy. Using novel connections to randomized Byzantine agreement lower bounds, we show that neither of these properties hold in general if $n \leq 3k+3t$.

We start by showing that we cannot handle all utilities variants if $n \leq 3k+3t$. Our proof exposes a new connection between utility variants and the problem of *Weak Byzantine Agreement* [24]. Lamport [24] showed that there is no deterministic protocol with bounded running time for *weak Byzantine agreement* if $t \geq n/3$. We prove a stronger lower bound for any randomized protocol that only assumes that the running time has finite expectation.

Proposition 1. *If* $\max\{2, k+t\} < n \leq 3k+3t$, *all* 2^n *input values are equally likely, and* P *is a (possibly randomized) protocol with finite expected running time (that is, for all protocols* P'' *and sets* $|T| \leq k+t$, *the expected running time of processes* P_{N-T} *given* (P_{N-T}, P''_T) *is finite), then there exists a protocol* P' *and a set* T *of players with* $|T| \leq k+t$ *such that an execution of* (P_{N-T}, P'_T) *is unsuccessful for the weak Byzantine agreement problem with nonzero probability.*

The idea of our impossibility result is to construct a game that captures weak Byzantine agreement. The challenge in the proof is that, while in the Byzantine agreement problem, nature chooses which processes are faulty, in the game, the players decide whether or not to behave in a faulty way. Thus, we must set up the incentives so that players gain by choosing to be faulty iff Byzantine agreement cannot be attained, while ensuring that a (k,t)-robust cheap-talk implementation of the mediator's strategy in the game will solve Byzantine agreement.

Theorem 3. *If $2k+2t < n \leq 3k+3t$, there is a game $\Gamma(u)$ and a strong (k,t)-robust equilibrium σ of a game Γ_d with a mediator d that extends Γ such that there exists a $(k+t)$-punishment strategy with respect to σ and there does not exist a strategy σ_{CT} such that for all utility variants $\Gamma(u')$ of $\Gamma(u)$, if σ is a (k,t)-robust equilibrium of $\Gamma_d(u')$, then $(\Gamma_{\mathrm{CT}}(u'),\sigma_{\mathrm{CT}})$ is a (k,t)-robust implementation of $(\Gamma_d(u'),\sigma)$.*

Theorem 3 shows that we cannot, in general, get a *uniform* implementation if $n \leq 3k+3t$. As shown in Theorem 1(b)–(e), we can implement mediators if $n \leq 3k+3t$ by taking advantage of knowing the players' utilities.

We next prove that if $2k+3t < n \leq 3k+3t$, although mediators can be implemented, they cannot be implemented without a punishment strategy. In fact we prove that they cannot even be ϵ–implemented without a punishment strategy. Barany [6] proves a weaker version of a special case of this result, where $n=3$, $k=1$, and $t=0$. It is not clear how to extend Barany's argument to the general case, or to ϵ–implementation. We use the power of randomized Byzantine agreement lower bounds for this result.

Theorem 4. *If $2k+2t < n \leq 3k+3t$, then there exists a game Γ, an $\epsilon > 0$, and a strong (k,t)-robust equilibrium σ of a game Γ_d with a mediator d that extends Γ, for which there does not exist a strategy σ_{CT} in the CT game that extends Γ such that σ_{CT} is an ϵ–(k,t)-robust implementation of σ.*

We now show that the assumption that $n > 2k+3t$ in Theorem 1 is necessary. More precisely, we show that if $n \leq 2k+3t$, then there is a game with a mediator that has a (k,t)-robust equilibrium that does not have a (k,t)-robust implementation in a cheap-talk game. We actually prove a stronger result: we show that there cannot even be an ϵ–(k,t)-robust implementation, for sufficiently small ϵ.

Theorem 5. *If $k+2t < n \leq 2k+3t$, there exists a game Γ, a strong (k,t)-robust equilibrium σ of a game Γ_d with a mediator d that extends Γ, a $(k+t)$-punishment strategy with respect to σ, and an $\epsilon > 0$, such that there does not exist a strategy σ_{CT} in the CT extension of Γ such that σ_{CT} is an ϵ–(k,t)-robust implementation of σ.*

The proof of Theorem 5 splits into two cases: (1) $2k+2t < n \leq 2k+3t$ and $t \geq 1$ and (2) $k+2t < n \leq 2k+2t$. For the first case, we use a reduction to a generalization of the Byzantine agreement problem called the (k,t)-Detect/Agree *problem*. This problem is closely related to the problem of *broadcast with extended consistency* introduced by Fitzi et al. [16].

Theorem 6. *If $2k+2t < n \leq 2k+3t$ and $t \geq 1$, there exists a game Γ, an $\epsilon > 0$, a strong (k,t)-robust equilibrium σ of a game Γ_d with a mediator d that extends Γ, and a $(k+t)$-punishment strategy with respect to σ, such that there does not exist a strategy σ_{CT} in the CT extension of Γ which is an ϵ–(k,t)-robust implementation of σ.*

We then consider the second case of Theorem 5, where $k + 2t < n \leq 2k + 2t$. Since we do not assume players know when other players have decided in the underlying game, our proof is a strengthening of the lower bounds of [30,22].

Theorem 7. *If $k+2t < n \leq 2k+2t$, there exist a game Γ, an $\epsilon > 0$, a mediator game Γ_d extending Γ, a strong (k,t)-robust equilibrium σ of Γ_d, and a $(k+t)$-punishment strategy ρ with respect to σ, such that there is no strategy σ_{CT} that is an ϵ–(k,t)-robust implementation of σ in the cheap-talk extension of Γ, even with broadcast channels.*

Our last lower bound using Byzantine agreement impossibility results gives a lower bound that matches the upper bound of Theorem 1(e) for the case that $n > k + 3t$. We show that a PKI cannot be set up on the fly if $n \leq k + 3t$. Our proof is based on a reduction to a lower bound for the (k,t)-*partial broadcast problem*, a novel variant of Byzantine agreement that can be viewed as capturing minimal conditions that still allow us to prove strong randomized lower bounds.

Theorem 8. *If $\max(2, k+t) < n \leq k + 3t$, then there is a game Γ, a strong (k,t)-robust equilibrium σ of a game Γ_d with a mediator d that extends Γ for which there does not exist a strategy σ_{CT} in the CT game that extends Γ such that σ_{CT} is an ϵ–(k,t)-robust implementation of σ even if players are computationally bounded and we assume cryptography.*

Tight Bounds on Running Time

We now turn our attention to running times. We provide tight bounds on the number of rounds needed to ϵ–implement equilibrium when $k+t < n \leq 2(k+t)$. When $2(k+t) < n$ then the expected running time is independent of the game utilities and independent of ϵ. We show that for $k + t < n \leq 2(k+t)$ this is not the case. The expected running time must depend on the utilities, and if punishment does not exist then the running time must also depend on ϵ.

Theorem 9. *If $k+t < n \leq 2(k+t)$ and $k \geq 1$, then there exists a game Γ, a mediator game Γ_d that extends Γ, a strategy σ in Γ_d, and a strategy ρ in Γ such that*

(a) for all ϵ and b, there exists a utility function $u^{b,\epsilon}$ such that σ is a (k,t)-robust equilibrium in $\Gamma_d(u^{b,\epsilon})$ for all b and ϵ, ρ is a (k,t)-punishment strategy with respect to σ in $\Gamma(u^{b,\epsilon})$ if $n > k + 2t$, and there does not exist an ϵ–(k,t)-robust implementation of σ that runs in expected time b in the cheap-talk extension $\Gamma_{\mathrm{CT}}(u^{b,\epsilon})$ of $\Gamma(u^{b,\epsilon})$;

(b) there exists a utility function u such that σ is a (k,t)-robust equilibrium in $\Gamma_d(u)$ and, for all b, there exists ϵ such that there does not exist an ϵ–(k,t)-robust implementation of σ^i that runs in expected time b in the cheap-talk extension $\Gamma_{\mathrm{CT}}(u)$ of $\Gamma(u)$.

This is true even if players are computationally bounded, we assume cryptography and there are broadcast channels.

Note that, in part (b), it is not assumed that there is a (k, t)-punishment strategy with respect to σ in $\Gamma(u)$. With a punishment strategy, for a fixed family of utility functions, we can implement an ϵ–(k, t)-robust strategy in the mediator game using cheap talk with running time that is independent of ϵ; with no punishment strategy, the running time depends on ϵ in general.

References

1. Abraham, I., Dolev, D., Gonen, R., Halpern, J.Y.: Distributed computing meets game theory: Robust mechanisms for rational secret sharing and multiparty computation. In: Proc. 25th ACM Symp. Principles of Distributed Computing, pp. 53–62 (2006)
2. Abraham, I., Dolev, D., Gonen, R., Halpern, J.Y.: Distributed computing meets game theory: Robust mechanisms for rational secret sharing and multiparty computation (unpublished manuscript, 2007)
3. Abraham, I., Dolev, D., Halpern, J.Y.: Lower bounds on implementing robust and resilient mediators. arXiv:0704.3646v2
4. 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)
5. Aumann, R.J.: Correlated equilibrium as an expression of Bayesian rationality. Econometrica 55, 1–18 (1987)
6. Barany, I.: Fair distribution protocols or how the players replace fortune. Mathematics of Operations Research 17, 327–340 (1992)
7. Ben-Or, M., Goldwasser, S., Wigderson, A.: Completeness theorems for non-cryptographic fault-tolerant distributed computation. In: Proc. 20th ACM Symp. Theory of Computing, pp. 1–10 (1988)
8. Ben-Porath, E.: Cheap talk in games with incomplete information. J. Economic Theory 108(1), 45–71 (2003)
9. Boneh, D., Naor, M.: Timed commitments. In: Bellare, M. (ed.) CRYPTO 2000. LNCS, vol. 1880, pp. 236–254. Springer, Heidelberg (2000)
10. Chaum, D., Crépeau, C., Damgard, I.: Multiparty unconditionally secure protocols. In: Proc. 20th ACM Symp. Theory of Computing, pp. 11–19 (1988)
11. Crawford, V.P., Sobel, J.: Strategic information transmission. Econometrica 50(6), 1431–1451 (1982)
12. Dodis, Y., Halevi, S., Rabin, T.: A cryptographic solution to a game theoretic problem. In: Bellare, M. (ed.) CRYPTO 2000. LNCS, vol. 1880, pp. 112–130. Springer, Heidelberg (2000)
13. Eliaz, K.: Fault-tolerant implementation. Review of Economic Studies 69(3), 589–610 (2002)
14. Even, S., Goldreich, O., Lempel, A.: A randomized protocol for signing contracts. Commun. ACM 28(6), 637–647 (1985)
15. Fischer, M.J., Lynch, N.A., Paterson, M.S.: Impossibility of distributed consensus with one faulty processor. Journal of the ACM 32(2), 374–382 (1985)
16. Fitzi, M., Hirt, M., Holenstein, T., Wullschleger, J.: Two-threshold broadcast and detectable multi-party computation. In: Biham, E. (ed.) EUROCRYPT 2003. LNCS, vol. 2656, pp. 51–67. Springer, Heidelberg (2003)
17. Forges, F.: Universal mechanisms. Econometrica 58(6), 1341–1364 (1990)

18. Goldreich, O.: Foundations of Cryptography, vol. 2. Cambridge University Press, Cambridge (2004)
19. Goldreich, O., Micali, S., Wigderson, A.: How to play any mental game. In: Proc. 19th ACM Symp. Theory of Computing, pp. 218–229 (1987)
20. Gordon, D., Katz, J.: Rational secret sharing, revisited. In: De Prisco, R., Yung, M. (eds.) SCN 2006. LNCS, vol. 4116, pp. 229–241. Springer, Heidelberg (2006)
21. Halpern, J.Y., Teague, V.: Rational secret sharing and multiparty computation: extended abstract. In: Proc. 36th ACM Symp. Theory of Computing, pp. 623–632 (2004)
22. Heller, Y.: A minority-proof cheap-talk protocol (2005) (Unpublished manuscript)
23. Izmalkov, S., Micali, S., Lepinski, M.: Rational secure computation and ideal mechanism design. In: Proc. 46th IEEE Symp. Foundations of Computer Science, pp. 585–595 (2005)
24. Lamport, L.: The weak byzantine generals problem. J. ACM 30(3), 668–676 (1983)
25. Lepinksi, M., Micali, S., Shelat, A.: Collusion-free protocols. In: Proc. 37th ACM Symp. Theory of Computing, pp. 543–552 (2005)
26. Lepinski, M., Micali, S., Peikert, C., Shelat, A.: Completely fair SFE and coalition-safe cheap talk. In: Proc. 23rd ACM Symp. Principles of Distributed Computing, pp. 1–10 (2004)
27. Lysyanskaya, A., Triandopoulos, N.: Rationality and Adversarial Behavior in Multi-party Computation. In: Dwork, C. (ed.) CRYPTO 2006. LNCS, vol. 4117, pp. 180–197. Springer, Heidelberg (2006)
28. Myerson, R.B.: Game Theory: Analysis of Conflict. Harvard University Press (September 1997)
29. Rabin, T., Ben-Or, M.: Verifiable secret sharing and multiparty protocols with honest majority. In: Proc. 21st ACM Symp. Theory of Computing, pp. 73–85 (1989)
30. Shamir, A., Rivest, R.L., Adelman, L.: Mental poker. In: Klarner, D.A. (ed.) The Mathematical Gardner, Prindle, Weber, Schmidt, Boston, Mass, pp. 37–43 (1981)
31. Urbano, A., Vila, J.E.: Computational complexity and communication: Coordination in two-player games. Econometrica 70(5), 1893–1927 (2002)
32. Urbano, A., Vila, J.E.: Computationally restricted unmediated talk under incomplete information. Economic Theory 23(2), 283–320 (2004)
33. Yao, A.: Protocols for secure computation (extended abstract). In: Proc. 23rd IEEE Symp. Foundations of Computer Science, pp. 160–164 (1982)