

Implementing Mediators with Asynchronous Cheap Talk

Ittai Abraham
ittai@cs.huji.ac.il
VMWARE

Danny Dolev
ittai@cs.huji.ac.il
The Hebrew University of Jerusalem

Ivan Geffner
ieg8@cornell.edu
Cornell University

Joseph Y. Halpern
halpern@cs.cornell.edu
Cornell University

October 20, 2019

Abstract

A mediator can help non-cooperative agents obtain an equilibrium that may otherwise not be possible. We study the ability of players to obtain the same equilibrium without a mediator, using only *cheap talk*, that is, nonbinding pre-play communication. Previous work has considered this problem in a synchronous setting. Here we consider the effect of asynchrony on the problem, and provide upper bounds for implementing mediators. Considering asynchronous environments introduces new subtleties, including exactly what solution concept is most appropriate and determining what move is played if the cheap talk goes on forever. Different results are obtained depending on whether the move after such “infinite play” is under the control of the players or part of the description of the game.

1 Introduction

Having a trusted mediator often makes solving a problem much easier. For example, a problem such as Byzantine agreement becomes trivial with a mediator: agents can just send their initial input to the mediator, and the mediator sends the majority value back to all the agents, which they then output. Not surprisingly, 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 has been on *secure multi-party computation* [Goldreich, Micali, and Wigderson 1987; Yao 1982]. Here it is assumed that each agent i has some private information x_i . Fix functions $f_1, \dots, f_n$. The goal is to have agent i learn $f_i(x_1, \dots, x_n)$ without learning anything about x_j for $j \neq i$ beyond what is revealed by the value of $f_i(x_1, \dots, x_n)$. 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 $f_i(x_1, \dots, x_n)$. Work on multiparty computation provides conditions under which this can be done in a synchronous system [Ben-Or, Goldwasser, and Wigderson 1988; Goldreich, Micali, and Wigderson 1987; Shamir, Rivest, and Adelman 1981; Yao 1982] and in an asynchronous system [Ben-Or, Canetti, and Goldreich 1993;

Ben-Or, Kelmer, and Rabin 1994]. 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.

In the computer science literature, the interest has been in performing multiparty computation in the presence of possibly malicious adversaries, who do everything they can to subvert the computation. In contrast, in the game theory literature, the assumption is that players have preferences and seek to maximize their utility; thus, they will subvert the computation iff it is in their best interests to do so. In [Abraham, Dolev, Gonen, and Halpern 2006; Abraham, Dolev, and Halpern 2008] (denoted ADGH and ADH, respectively, in the rest of the paper), it was argued that it is important to consider deviations by both rational players, who have preferences and try to maximize them, and players that we can view as malicious, although it is perhaps better to think of them as rational players whose utilities are not known by the mechanism designer (or other players). ADGH and ADH 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. Tight bounds were proved on the ability to implement a (k, t) -robust equilibrium in the game with a mediator using cheap talk in synchronous systems. These bounds depend on, among other things, (a) the relationship between k , t and n , the total number of players in the system; (b) whether players know the exact utilities of the rational players; and (c) whether the game has a *punishment strategy*, where an m -punishment strategy is a strategy profile that, if used by all but at most m players, guarantees that every player gets a worse outcome than they do with the equilibrium strategy. The following is a high-level overview of results proved in the synchronous setting that will be of most relevance here. For these results, we assume that the communication with the mediator is bounded, it lasts for at most N rounds, and that the mediator can be represented by an arithmetic circuit of depth c .

- R1. If $n > 3k + 3t$, then a mediator can be implemented using cheap talk; no punishment strategy is required, no knowledge of other agents' utilities is required, and the cheap-talk protocol has bounded running time $O(nNc)$, independent of the utilities.
- R2. If $n > 2k + 3t$, then a mediator can be implemented using cheap talk if there is a $(k + t)$ -punishment strategy and the utilities of the rational players are known; the cheap-talk protocol has expected running time $O(nNc)$. (In R2, unlike R1, the cheap-talk protocol may be unbounded, although it has finite expected running time.)

In ADH, lower bounds are presented that match the upper bounds above. Thus, for example, it is shown that $n > 3k + 3t$ is necessary in R1; if $n \leq 3k + 3t$, then we cannot implement a mediator in general if we do not have a punishment strategy or if the utilities are unknown. The proofs of the upper bounds make heavy use of the fact that the setting is synchronous. Here we consider the impact of asynchrony on these results. Once we introduce asynchrony, we must revisit the question of what it even means to implement an equilibrium using cheap talk. Notions like (Bayesian) Nash equilibrium implicitly assume that all uncertainty is described probabilistically. Having a probability is necessary to talk about an agent's expected utility, given that a certain strategy profile is played. If we were willing to put a distribution on how long messages take to arrive and on when agents are scheduled to move, then we could apply notions like Nash equilibrium without difficulty. However, it is notoriously difficult to quantify this uncertainty. The typical approach used to analyze algorithms in the presence of uncertainty that is not quantified probabilistically is

to assume that all the non-probabilistic uncertainty is resolved by the environment according to some strategy. Thus, the environment uses some strategy to decide when each agent will be allowed to play and how long each message takes to be delivered. The algorithm is then proved correct no matter what strategy the environment is following in some class of strategies. For example, we might restrict the environment’s strategy to being *fair*, so that every agent eventually gets a chance to move. (See [Halpern and Tuttle 1993] for a discussion of this approach and further references.)

We follow this approach in the context of games. Note that once we fix the environment’s strategy, we have an ordinary game, where uncertainty is quantified by probability. In this setting, we consider *ex post equilibrium*. A strategy is an *ex post equilibrium* if it is an equilibrium no matter what strategy the environment uses. *Ex post equilibrium* is a strong notion, but, as we show by example, it can often be attained with the help of a mediator. It is arguably the closest analogue to Nash equilibrium in an asynchronous setting.

Another issue that plays a major role in an asynchronous setting is what happens if the strategies of players result in some players being *livelocked*, talking indefinitely without making a move in the underlying game, or in some players being *deadlocked*, waiting indefinitely without moving in the underlying game. We consider two approaches for dealing with this problem. One is called the *default-move approach*. In this approach, as part of the description of the game, there is a default move for each player which is imposed if that player fails to explicitly make a move in the cheap-talk phase. Aumann and Hart [2003] considered a different approach, which we henceforth call the *AH approach*, where a player’s strategy in the underlying game is a function of the (possibly infinite) history of the player in the cheap-talk phase. We can think of this almost as if the player writes a will, describing what he would like to have done (as a function of the history) if the game ends before he has had a chance to move.

We believe that both the AH approach and the default-move approach are reasonable in different contexts. The AH approach makes sense if the agent can leave instructions that will be carried out by an “executor” if the cheap-talk game deadlocks. But if we consider a game-theoretic variant of Byzantine agreement, it seems more reasonable to say that if a malicious agent can prevent an agent from making a move in finite time, the agent should not get a chance to make a move after the cheap-talk phase has ended.

Our results show that, in the worst case, the cost of asynchrony is an extra $k + t$ in the bounds on n , but we can sometimes save k or even $k + t$ if there is a punishment strategy or if we are willing to tolerate an ϵ “error”. For example, with both the AH approach and the default-move approach, if the utilities are not known, we can implement a mediator using asynchronous cheap talk if $n > 4k + 4t$. Thus, compared to R1, we need an extra $k + t$. However, if we are willing to accept a small probability of error, so that rather than implementing the mediator we get only an ϵ -implementation, and are also willing to accept ϵ -(k, t)-robustness (which, roughly speaking, means that players get within ϵ of the best they could get), then we can do this if $n > 3k + 3t$, again, using both the AH approach and the default-move approach.

Just as in the synchronous case, we can do better if we assume that there is a punishment strategy and utilities are known (as in R2). Specifically, with the AH approach, we can implement a mediator if $n > 3k + 4t$ (compared to $n > 2k + 3t$ in the synchronous case), and can ϵ -implement a mediator if $n > 2k + 3t$. We use the punishment to deal with deadlock. If a good player is waiting for a message that never arrives, then the waiting player instructs his executor to carry out a punishment in his will. Having a punishment does not seem to help in the default-move approach unless the

default move is a punishment; if it is, then we can get the same results as with the AH approach.

If there is a punishment strategy, these results significantly improve those of Even, Goldreich, and Lempel [1985]. They provide a protocol with similar properties, but the expected number of messages sent is $O(1/\epsilon)$; with a punishment strategy, we show that a bounded number of messages can be sent, with the bound being independent of ϵ .

The rest of this paper is organized as follows. In Section 2, we review all the relevant definitions. In Section 3, we review the definitions of the solution concepts from ADGH that we use. In Section 4 we state our results carefully. In Section 5, we discuss the security-theoretic notion of *t-bisimulation* needed in our proofs, and state results from a companion paper [Geffner and Halpern 2018] regarding *t-bisimulation*. We outline the proofs of our main theorems in Section 6, and conclude in Section 7.

2 Definitions

Asynchronous games, mediator games, and cheap talk: 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 an extension Γ_{CT} of Γ with (asynchronous) cheap talk. We assume that Γ is a *normal-form Bayesian game*: each player has a type t taken from some type space $\mathcal{T}_i$, such that there is a commonly known distribution on $\mathcal{T} \subseteq \mathcal{T}_1 \times \dots \times \mathcal{T}_n$, the set of types; each player i chooses an action $a \in A_i$, the set of actions of agent i ; player i 's utility u_i is determined by the type profile of the players and the actions they take. A strategy for player i in the Bayesian game is just a function T_i to A_i , which tells player i what to do, given his type. If $A = A_1 \times \dots \times A_n$, then a strategy profile $\vec{\sigma} = (\sigma_1, \dots, \sigma_n)$ can be viewed as a function $\vec{\sigma} : \mathcal{T} \rightarrow \Delta(A)$ (where, as usual, $\Delta(X)$ denotes the set of probability distributions on X).

The basic notions of a game with a mediator, a cheap-talk game, and implementation are standard in the game-theory literature. However, since we consider them in an asynchronous setting, we must modify the definitions somewhat.

We first define *asynchronous games*. In an asynchronous game, we assume that players alternate making moves with the environment—first the environment moves, then a player moves, then the environment moves, and so on. The environment's move consists of choosing a player i to move next and a set of messages in transit to i that will be delivered just before i moves (so that i 's move can depend on the messages i receives). The environment is subject to two constraints: all messages sent must eventually be delivered and, for all times m and players i , if i is still playing the game at time m , then there must be some time $m' \geq m$ that i is chosen to move. We can describe an asynchronous game by a game tree. Associated with each non-leaf node or history is either a player—the player whose move it is at that node—or the environment (note that both the players and the environment can use probabilistic strategies). The nodes where a player i moves are further partitioned into *information sets*; intuitively, these are nodes that player i cannot tell apart. We assume that the environment has complete information, so that the environment's information sets just consist of the singletons. A *strategy* for player i is a (possibly randomized) function from i 's information sets to actions; we can similarly define a strategy for the environment. We can essentially view the environment strategy as defining a scheduler (and thus we sometimes refer to an environment strategy as a scheduler).

For our results, we start with an n -player Bayesian game Γ in normal form (called the *underlying game*), with $\{1, \dots, n\}$ being the set of players, and then consider two games that *extend* Γ . A game Γ' extends Γ if the players have initial types from the same type space as Γ , with the same distribution over types; moreover, in each path of the game tree for Γ' , the players send and receive messages, and perform at most one action from Γ . In a history where each player makes a move from Γ , each player gets the same utility as in Γ (where the utility is a function of the moves made and the types). That leaves open the question of what happens in a complete history of Γ' where some players do not make a move in Γ . As we suggested in the introduction, we consider two approaches to dealing with this. In the first approach, we assume that the description of Γ' includes a function M_i for each player i that maps player i 's type to a move in Γ . In an infinite history h where i has type t and does not make a move in Γ , i is viewed as having made move $M_i(t)$. We can then define each player's utility in h as above. This is the *default-move approach*. In the AH approach, we extend the notion of strategy so that i 's strategy in Γ' also describes what move i makes in the underlying game Γ in any infinite history h where i has not made a move in Γ . In the AH approach, i 's move in h is under i 's control; in the default-move approach, it is not.

Given an underlying Bayesian game Γ (which we assume is synchronous—the players move simultaneously), we will be interested in two types of extensions. A *mediator game* extending Γ is an asynchronous game where players can send messages to and receive messages from a mediator (who can be viewed as a trusted third party) as well as making a move in Γ ; “good” or “honest” players do not send messages to each other, but “bad” players (i.e., one of the k rational deviating players or one of the t “malicious” players with unknown utilities) may send messages to each other as well as to the mediator. We assume that the space of possible messages that can be sent in a mediator game is fixed and finite.

In an asynchronous cheap-talk game extending Γ , there is no mediator. Players send messages to each other via asynchronous channels, as well as making a move in Γ . We assume that each pair of agents communicates over an *asynchronous authenticated private channel*, so the adversary cannot eavesdrop on conversations between the players, and players can identify the sender of each message. Finally, we assume that in both the mediator game and the cheap-talk game, when a player is first scheduled, it gets a signal that the game has started (either an external signal from the environment, or a game-related message from another player or the mediator).

Implementation: In the synchronous setting, a strategy profile $\vec{\sigma}'$ in a cheap-talk game Γ_{CT} extending an underlying game Γ *implements* a strategy $\vec{\sigma}$ in a mediator game Γ_d extending Γ if $\vec{\sigma}$ and $\vec{\sigma}'$ correspond to the same strategy in Γ ; that is, they induce the same function from $\mathcal{T}$ to $\Delta(A)$. The notion of implementation is more complicated in an asynchronous setting, because the probability on action profiles also depends on the environment strategy. Because Γ_{CT} and Γ_d are quite different games, the environment's strategies in Γ_{CT} are quite different from those in Γ_d . So we now say that $\vec{\sigma}'$ implements $\vec{\sigma}$ if the set of distributions on actions profiles in Γ induced by $\vec{\sigma}$ and all possible choices of environment strategy is the same as that induced by $\vec{\sigma}'$ and all possible choices of environment strategy. More precisely, let $\mathcal{S}_{\Gamma',e}$ and $\mathcal{S}_{\Gamma'',e}$ denote the set of environment strategies in Γ' and Γ'' , respectively. A strategy $\sigma_e \in \mathcal{S}_{\Gamma',e}$ and a strategy profile $\vec{\sigma}$ for the players in Γ' together induce a function $(\vec{\sigma}, \sigma_e)$ from $\mathcal{T}$ to $\Delta(A)$. A strategy profile $\vec{\sigma}'$ in Γ' *implements* a strategy profile $\vec{\sigma}''$ in Γ'' if $\{(\vec{\sigma}', \sigma'_e) : \sigma'_e \in \mathcal{S}_{\Gamma',e}\} = \{(\vec{\sigma}'', \sigma''_e) : \sigma''_e \in \mathcal{S}_{\Gamma'',e}\}$. Since the outcome that arises if the players use a particular strategy may depend on what the environment does, this says that the set of outcomes that can result if the players use $\vec{\sigma}'$ is the same as the set of outcomes that can result if the players use $\vec{\sigma}''$.

For some of our results, we cannot get an exact implementation; there may be some error. Given two discrete distributions π and π' on some space S , the *distance* between π and π' , denoted $\text{dist}(\pi, \pi')$, is at most ϵ if $\sum_{s \in S} |\pi(s) - \pi'(s)| \leq \epsilon$. As we observed earlier, in the mediator game and the cheap-talk game, a strategy profile $\vec{\sigma}$ for the players and a strategy σ_e for the environment together induce a mapping from type profiles to $\Delta(A)$. We lift the notion of distance to such function by defining $\text{dist}((\vec{\sigma}, \sigma_e), (\vec{\sigma}', \sigma'_e)) = \max_{\vec{x} \in \mathcal{T}} \text{dist}((\vec{\sigma}, \sigma_e)(\vec{x}), (\vec{\sigma}', \sigma'_e)(\vec{x}))$. Say that $\vec{\sigma}'$ ϵ -implements $\vec{\sigma}''$ if

- for all $\sigma'_e \in \mathcal{S}_{\Gamma', e}$ there exists $\sigma''_e \in \mathcal{S}_{\Gamma'', e}$ such that $\text{dist}((\vec{\sigma}', \sigma'_e), (\vec{\sigma}'', \sigma''_e)) \leq \epsilon$; and
- for all $\sigma''_e \in \mathcal{S}_{\Gamma'', e}$ there exists $\sigma'_e \in \mathcal{S}_{\Gamma', e}$ such that $\text{dist}((\vec{\sigma}'', \sigma''_e), (\vec{\sigma}', \sigma'_e)) \leq \epsilon$.

Note that $\vec{\sigma}'$ implements $\vec{\sigma}''$ iff $\vec{\sigma}'$ 0-implements $\vec{\sigma}''$.

The notion of implementation is quite strong. For example, if $\vec{\sigma}'$ involves fewer rounds of communication than $\vec{\sigma}''$, there may be far fewer distinct schedulers in the game involving $\vec{\sigma}'$ than in the game involving $\vec{\sigma}''$. Thus, we may not be able to recover the effect of all possible schedulers. (Indeed, for some of our results the implementation needs to be quite long precisely in order to capture all possible schedulers.) This suggests the following notion: a strategy profile $\vec{\sigma}'$ in Γ' *weakly implements* a strategy profile $\vec{\sigma}''$ in Γ'' if $\{(\vec{\sigma}', \sigma'_e) : \sigma'_e \in \mathcal{S}_{\Gamma', e}\} \subseteq \{(\vec{\sigma}'', \sigma''_e) : \sigma''_e \in \mathcal{S}_{\Gamma'', e}\}$. Thus, if $\vec{\sigma}'$ weakly implements $\vec{\sigma}''$, then every outcome of $\vec{\sigma}'$ is one that could also have arisen with $\vec{\sigma}''$, but the converse may not be true. Specifically, there may be some behaviors of the environment with $\vec{\sigma}''$ that cannot be simulated by $\vec{\sigma}'$. As we shall see, this may actually be a feature: we can sometimes simulate the effect of only “good” schedulers. In any case, note that in the synchronous setting, implementation and weak implementation coincide. We can also define a notion of weak ϵ -implementation in the obvious way; we leave the details to the reader.

Termination: We will be interested in asynchronous games where, almost surely, the honest players stop sending messages and make a move in the underlying game. In the mediator games that we consider, this happens after only a bounded number of messages have been sent. But even with this bound, there may not be a point in a history when players know that they can stop sending messages; although a player i may have moved in the underlying game, i may still need to keep checking for incoming message, and may need to respond to them, to ensure that other players can make the appropriate move.

For some of our results, we must assume that, in the mediator game, there comes a point when all honest players know that they have terminated the protocol; they will not get further messages from the mediator and can stop sending messages to the mediator, and should make a move in the underlying game if they have not done so yet. For simplicity, for these results, we restrict the honest players and the mediator to using strategy profiles that have the following *canonical form*: Using a canonical strategy, player i may send a message to the mediator whenever it is scheduled, as long as i has not received a message from the mediator containing “STOP”. If player i gets a message from the mediator that includes “STOP”, then i makes a move in the underlying game and halts. We assume that, as long as the honest players and mediator follow their part of the canonical strategy profile, there is a constant r such that, no matter what strategy the rational and malicious players and the environment use, the mediator sends each player i at most r messages in each history, and the final message includes “STOP”. We conjecture that the assumption that players and mediator are using a strategy in canonical form in the mediator game is without loss

of generality; that is, a (k, t) -robust strategy profile in a mediator game Γ_d can be implemented by a (k, t) -robust strategy profile in Γ_d that is in canonical form.

3 Solution concepts

In this section, we review the solution concepts introduced in ADGH and extend them to asynchronous settings.

Note that in an asynchronous game Γ , the utility of a player i can depend not only on the strategies of the agents, but on what the environment does. Since we consider an underlying game, a mediator game, and a cheap-talk game, it is useful to include explicitly in the utility function which game is being considered. Thus, we write $u_i(\Gamma, \vec{\sigma}, \sigma_d, \sigma_e, \vec{x})$ to denote the expected utility of player i in game Γ when players play strategy profile $\vec{\sigma}$, the mediator plays σ_d , the environment plays σ_e , and the type profile is $\vec{x}$. We typically say “input profile” rather than “type profile”, since in our setting, the type of player i is just i ’s initial input. Note that if Γ is the underlying game, the σ_e component is unnecessary, since the underlying game is assumed to be synchronous. We occasionally omit the mediator strategy σ_d when it is clear from context.

Given a type space $\mathcal{T}$, a set K of players, and $\vec{x} \in \mathcal{T}$, let $\mathcal{T}(\vec{x}_K) = \{\vec{x}' : \vec{x}'_K = \vec{x}_K\}$. If Γ is a Bayesian game over type space $\mathcal{T}$, $\vec{\sigma}$ is a strategy profile in Γ , and $\Pr$ is the probability on the type space $\mathcal{T}$, let

$$u_i(\Gamma, \vec{\sigma}, \sigma_e, \vec{x}_K) = \sum_{\vec{x}' \in \mathcal{T}(\vec{x}_K)} \Pr(\vec{x}' \mid \mathcal{T}(\vec{x}_K)) u_i(\Gamma, \vec{\sigma}, \sigma_e, \vec{x}').$$

Thus, $u_i(\Gamma, \vec{\sigma}, \sigma_e, \vec{x}_K)$ is i ’s expected payoff if everyone uses strategy $\vec{\sigma}$ and type profiles are in $\mathcal{T}(\vec{x}_K)$.

k -resilient equilibrium: In a standard game, a strategy profile is a Nash equilibrium if no player can gain any advantage by using a different strategy, given that all the other players do not change their strategies. The notion of k -resilient equilibrium extends Nash equilibrium to allow for coalitions.

Definition 3.1. $\vec{\sigma}$ is a k -resilient equilibrium (resp., strongly k -resilient equilibrium) in an asynchronous game Γ if, for all subsets K of players with $1 \leq |K| \leq k$, all strategy profiles $\vec{\tau}_K$ for the players in K , all type profiles $\vec{x} \in \mathcal{T}$, and all strategies σ_e of the environment, $u_i(\Gamma, (\vec{\sigma}_{-K}, \vec{\tau}_K), \sigma_e, \vec{x}_K) \leq u_i(\Gamma, \vec{\sigma}, \sigma_e, \vec{x}_K)$ for some (resp., all) $i \in K$.¹

Thus, $\vec{\sigma}$ is k -resilient if, no matter what the environment does, no subset K of at most k players can all do better by deviating, even if they share their type information (so that if the true type is $\vec{x}$, the players in K know $\vec{x}_K$). It is strongly k -resilient if not even one of the players in K can do better if all the players in K deviate.

For some of our results we will be interested in equilibria that are “almost” k -resilient, in the sense that no player in a coalition can do more than ϵ better if the coalition deviates from the protocol, for some small ϵ .

¹As usual, in the strategy profile $(\vec{\sigma}_{-K}, \vec{\tau}_K)$, each player $i \in K$ plays τ_i and each player $i \notin K$ plays σ_i .

Definition 3.2. For $\epsilon > 0$, $\vec{\sigma}$ is an ϵ - k -resilient equilibrium (resp., strongly ϵ - k -resilient equilibrium) if, for all subsets K of players, all strategy profiles $\vec{\tau}_K$ for the players in K , all type profiles $\vec{x} \in \mathcal{T}$, and all strategies σ_e of the environment, we have $u_i(\Gamma, (\vec{\sigma}_{-K}, \vec{\tau}_K), \sigma_e, \vec{x}_K) < u_i(\Gamma, \vec{\sigma}, \sigma_e, \vec{x}_K) + \epsilon$ for some (resp., for all) $i \in K$.

Note that we have “ $< u_i(\Gamma, \vec{\sigma}, \sigma_e, \vec{x}_K) + \epsilon$ ” here, not “ $\leq$ ”; this means that a 0- k -resilient equilibrium is not a k -resilient equilibrium. However, an equilibrium is k -resilient iff it is ϵ - k -resilient for all $\epsilon > 0$. We have used this slightly nonstandard definition to make the statements of our theorems cleaner.

Robustness: A standard assumption in game theory is that utilities are (commonly) known; when we are given a game we are also given each player’s utility. When players make decision, they can take other players’ utilities into account. However, in large systems, it seems almost invariably the case that there will be some fraction of users who do not respond to incentives the way we expect. For example, in a peer-to-peer network like Kazaa or Gnutella, it would seem that no rational agent should share files. Whether or not you can get a file depends only on whether other people share files; on the other hand, it seems that there are disincentives for sharing (the possibility of lawsuits, use of bandwidth, etc.). Nevertheless, people do share files. However, studies of the Gnutella network have shown almost 70 percent of users share no files and nearly 50 percent of responses are from the top 1 percent of sharing hosts [Adar and Huberman 2000]. It seems important to design protocols that tolerate such unanticipated behaviors, so that the payoffs of the users who follow the recommended strategy are not affected by players who deviate, provided that not too many deviate.

Definition 3.3. A strategy profile $\vec{\sigma}$ is t -immune in a game Γ if, for all subsets T of players with $|T| \leq t$, all strategy profiles $\vec{\tau}$, all $i \notin T$, all type profiles $\vec{x} \in \mathcal{T}$, and all strategies σ_e of the environment, we have $u_i(\Gamma, (\vec{\sigma}_{-T}, \vec{\tau}_T), \sigma_e, \vec{x}_T) \geq u_i(\Gamma, \vec{\sigma}, \sigma_e, \vec{x}_T)$.

Intuitively, $\vec{\sigma}$ is t -immune if there is nothing that players in a set T of size at most t can do to give the players not in T a worse payoff, even if the players in T share their type information.

The notion of t -immunity and k -resilience address different concerns. For t -immunity, we consider the payoffs of the players not in K ; for k -resilience, we consider the payoffs of players in K . It is natural to combine both notions. Given a strategy profile $\vec{\tau}$, let $\Gamma_{\vec{\tau}}^T$ be the game which is identical to Γ except that the players in T are fixed to playing strategy $\vec{\tau}_T$.

Definition 3.4. $\vec{\sigma}$ is a (strongly) (k, t) -robust equilibrium in a game Γ if $\vec{\sigma}$ is t -immune and, for all subsets T of players with $|T| \leq t$ and all strategy profiles $\vec{\tau}$, $(\vec{\sigma}_{-T}, \vec{\tau}_T)$ is a (strongly) k -resilient equilibrium of $\Gamma_{\vec{\tau}}^T$.

We can define “approximate” notions of t -immunity and (k, t) -robustness analogous to Definition 3.2:

Definition 3.5. For $\epsilon > 0$, a strategy profile $\vec{\sigma}$ is ϵ - t -immune in Γ if, for all subsets T of players with $|T| \leq t$, all strategy profiles $\vec{\tau}$, all $i \notin T$, all type profiles $\vec{x} \in \mathcal{T}$, and all strategies σ_e of the environment, we have $u_i(\Gamma, (\vec{\sigma}_{-T}, \vec{\tau}_T), \sigma_e, \vec{x}_T) > u_i(\Gamma, \vec{\sigma}, \sigma_e, \vec{x}_T) - \epsilon$.

Definition 3.6. For $\epsilon \geq 0$, $\vec{\sigma}$ is a (strongly) ϵ -(k, t)-robust equilibrium in Γ if $\vec{\sigma}$ is ϵ - t -immune and, for all subsets T of players with $|T| \leq t$ and strategy profiles $\vec{\tau}_T$, $(\vec{\sigma}_{-T}, \vec{\tau}_T)$ is a (strongly) ϵ - k -resilient equilibrium of $\Gamma_{\vec{\tau}}^T$.

4 Main theorems: formal statements

In this section, we state our results formally. We begin with a result that is an analogue of R1 in the asynchronous setting. We say that a game Γ' is a utility variant of a game Γ if Γ' and Γ have the same game tree, but the utilities of the players may be different in Γ and Γ' . We use the notation $\Gamma(\vec{u})$ if we want to emphasize that $\vec{u}$ is the utility function in game Γ . We then take $\Gamma(\vec{u}')$ to be the utility variant of Γ with utility function $\vec{u}'$.

Two more technical comments before stating the theorems: in the mediator game we can also view the mediator as a player (albeit one without a utility function) that is following a strategy. Thus, when we talk about a strategy profile that is a (k, t) -robust equilibrium in the mediator game, we must give the mediator's strategy as well as the players' strategies. We sometimes write $\vec{\sigma} + \sigma_d$ if we want to distinguish the players' strategy profile $\vec{\sigma}$ from the mediator's strategy σ_d . We occasionally abuse notation and drop the σ_d if it is clear from context, and just talk about $\vec{\sigma}$ being a (k, t) -robust equilibrium.

In general, we do not have a bound on the number of messages sent in our implementation. However, we can get a polynomial bound if we restrict the mediator to using what we call a *responsive* strategy. Roughly speaking, a responsive strategy $\vec{\sigma} + \sigma_d$ is one where (1) the mediator sends messages only initially or in response to receiving a message since the last time it moved; (2) if it has received a message since the last time it was scheduled, it acts just as it would have acted if the message had been received immediately after the last time it acted (rather than possibly being received after the mediator was scheduled a number of times without receiving a new message); and (3) there is a bound C such that the mediator uses at most C random bits (which are set to 1 with probability $1/2$) in each run; in this case we say that $\vec{\sigma} + \sigma_d$ has *uniformly bounded randomization*. See [Geffner and Halpern 2018] for the formal definition.

Theorem 4.1. *If Γ is a normal-form Bayesian game with n players, $\vec{\sigma} + \sigma_d$ is a strategy profile for the players and the mediator in an asynchronous mediator game Γ_d that extends Γ , and $n > 4k + 4t$, then with both the default-move approach and the AH approach, there exists a strategy profile $\vec{\sigma}_{CT}$ that implements $\vec{\sigma} + \sigma_d$ in the asynchronous cheap-talk game Γ_{CT} such that for all utility variants $\Gamma_d(\vec{u}')$ of Γ_d , if $\vec{\sigma} + \sigma_d$ is a (strongly) (k, t) -robust equilibrium in $\Gamma_d(\vec{u}')$, then $\vec{\sigma}_{CT}$ is a (strongly) (k, t) -robust equilibrium in $\Gamma_{CT}(\vec{u}')$. If σ_d is responsive, then the number of messages sent in a history of $\vec{\sigma}_{CT}$ is polynomial in n and N , linear in c , and independent of $\vec{u}'$.*

The proof of Theorem 4.1 uses ideas from the multiparty computation protocol of Ben-Or, Canetti, and Goldreich [1993] (BCG from now on). Our construction actually needs stronger properties than these provided by BCG; we show that we can get protocols with these stronger properties in a companion paper [Geffner and Halpern 2018]; see Section 5 for further discussion.

We can obtain better bounds if we are willing to accept ϵ -equilibrium, using ideas due to Ben-Or, Kelmer, and Rabin [1994].

Theorem 4.2. *If Γ is a normal-form Bayesian game with n players, $\vec{\sigma} + \sigma_d$ is a strategy profile for the players and mediator in an asynchronous mediator game Γ_d that extends Γ , $M > 0$, and $n > 3k + 3t$, then with both the default-move approach and the AH approach, for all $\epsilon > 0$, there exists a strategy profile $\vec{\sigma}_{CT}$ in the asynchronous cheap-talk game Γ_{CT} that ϵ -implements $\vec{\sigma}$ such that for all utility variants $\Gamma_d(\vec{u}')$ of Γ_d bounded by $M/2$ (i.e., where the range of u'_i is contained in*

$[-M/2, M/2]$), if $\vec{\sigma} + \sigma_d$ is a (strongly) (k, t) -robust equilibrium in $\Gamma_d(\vec{u}')$, then $\vec{\sigma}_{CT}$ is a (strongly) ϵ -(k, t)-robust equilibrium in $\Gamma_{CT}(\vec{u}')$. If σ_d is responsive, then the number of messages sent in a history of $\vec{\sigma}_{CT}$ is polynomial in n and N , linear in c , and independent of $\vec{u}'$.

If we have a punishment strategy and utilities are known, we can do better with the AH approach. To make this precise, we need the definition of an m -punishment strategy [Abraham, Dolev, Gonen, and Halpern 2006] (which generalizes the notion of punishment strategy defined by Ben Porath [2003]). Before defining this carefully, note that in an asynchronous setting (i.e., in the mediator game and the cheap-talk game, but not in the underlying game), the utility of players depends on the environment's strategy as well as the players' strategy profile and the players' type profile.

Definition 4.3. If Γ' is an extension of an underlying game Γ , a strategy profile $\vec{\rho}$ in Γ is a k -punishment strategy with respect to a strategy profile $\vec{\sigma}'$ in Γ' if for all subsets K of players with $1 \leq |K| \leq k$, all strategy profiles $\vec{\sigma}$ in Γ , all strategies σ_e for the environment, all type profiles $\vec{x} \in \mathcal{T}$, and all players $i \in K$, we have

$$u_i(\Gamma', \vec{\sigma}', \sigma_e, \vec{x}_K) > u_i(\Gamma, (\vec{\sigma}_K, \vec{\rho}_{-K}), \vec{x}_K).$$

Thus, if $\vec{\rho}$ is a k -punishment strategy with respect to $\vec{\sigma}'$ and all but k players play their part of $\vec{\rho}$ in the underlying game, then all of the remaining players will be worse off than they would be in Γ' if everyone had played $\vec{\sigma}'$, no matter what they do in the underlying game.

Theorem 4.4. If Γ is a normal-form Bayesian game with n players, $\vec{\sigma} + \sigma_d$ is a strategy profile in canonical form with uniformly bounded randomization for the players and mediator in an asynchronous mediator game Γ_d that extends Γ , there is a $(k + t)$ -punishment strategy with respect to $\vec{\sigma} + \sigma_d$, and $n > 3k + 4t$, then with the AH approach, there exists a strategy profile $\vec{\sigma}_{CT}$ that implements $\vec{\sigma} + \sigma_d$ in the asynchronous cheap-talk game Γ_{CT} , and if $\vec{\sigma} + \sigma_d$ is a (strongly) (k, t) -robust equilibrium in Γ_d , then $\vec{\sigma}_{CT}$ is a (strongly) (k, t) -robust equilibrium in Γ_{CT} . If we require only that $\vec{\sigma}_{CT}$ is a weak implementation, then the number of messages in a history of $\vec{\sigma}_{CT}$ is polynomial in n and linear in c .

Note that in Theorem 4.4, the running time of the algorithm is significantly affected by whether we want $\vec{\sigma}_{CT}$ to implement $\vec{\sigma}$ or whether a weak implementation suffices.

If we assume both that there is a $(2k + 2t)$ -punishment strategy and that utilities are known, we can get an analogue to R2, but with an ϵ error.

Theorem 4.5. If Γ is a normal-form Bayesian game with n players, $\vec{\sigma} + \sigma_d$ is a strategy profile in canonical form with uniformly bounded randomization for the players and mediator in an asynchronous mediator game Γ_d that extends Γ , there is a $(2k + 2t)$ -punishment strategy with respect to $\vec{\sigma} + \sigma_d$, and $n > 2k + 3t$, then with the AH approach, for all $\epsilon > 0$ there is a strategy $\vec{\sigma}_{CT}$ that ϵ -implements $\vec{\sigma}$ in the asynchronous cheap-talk game Γ_{CT} such that if $\vec{\sigma} + \sigma_d$ is a (strongly) (k, t) -robust equilibrium in Γ_d , then $\vec{\sigma}_{CT}$ is a (strongly) ϵ -(k, t)-robust equilibrium in Γ_{CT} . If we require only that $\vec{\sigma}_{CT}$ is a weak implementation, then the number of messages in a history of $\vec{\sigma}_{CT}$ is polynomial in n and linear in c .

We prove these results using ideas in the spirit of ADGH, but much more care must be taken to deal with asynchrony. Among other things, we need stronger security guarantees than are traditionally provided for multiparty communication; see Section 5 for details.

5 Bisimulation and cotermination

Ben-Or, Goldwasser, and Wigderson [1988] (BGW from now on) and Ben-Or, Canetti, and Goldreich [1993] (BCG from now on) show that if a function f of n inputs provided by n players can be computed using a mediator, then it can be computed by the players without the mediator and without revealing any information beyond the function value, even when some of the players are malicious. BGW deal with the synchronous case and provide a protocol that tolerates up to $n/3$ malicious players; BCG deal with the asynchronous case and provide a protocol that tolerates up to $n/4$. The notion of not revealing any information is made precise by defining a set of *ideal* distributions over possible values of the function, and ensuring that the *real* distribution is identically distributed to one of those (see BGW and BCG for formal definitions and details).

We can view a mediator game as computing an action profile in the underlying game; the ideal distributions are the possible distributions over action profiles when the honest players play their component of the (k, t) -robust equilibrium strategy profile in the mediator game. BCG's protocol then essentially gives us a strategy in the cheap-talk game. However, the BCG protocol is not sufficient for our purposes for two reasons: it does not guarantee that the real protocol is an implementation of the ideal protocol in the sense of the definition in Section 2 (although it does suffice for *weak* implementation), nor does it guarantee that the protocol is a (k, t) -robust equilibrium. To prove these stronger results, we show that $\vec{\sigma}_{CT}$ can be constructed so as to satisfy some additional security properties, which we now define. For some of these definitions, it will be useful to introduce the notion of *adversary*. An adversary A is a tuple $(T, \vec{\tau}_T, \sigma_e)$ consisting of a set T of malicious players, a strategy $\vec{\tau}_T$ for players in T , and a strategy σ_e for the scheduler. To keep the notation as simple as possible, we omit the set T when it is clear from context, just writing A as the pair $(\vec{\tau}_T, \sigma_e)$. (Note that $\vec{\tau}_T$ is just a tuple of $|T|$ strategies. We write this tuple as $\vec{\tau}_T$ only because we view it as the strategies played by the players in T in the profile $\vec{\tau}$.)

Definition 5.1 (*t*-bisimulation). *Let $\vec{\pi}(\vec{x}, A)$ be the distribution over outputs when running strategy $\vec{\pi}$ with adversary $A = (\tau_T, \sigma_e)$. Protocol $\vec{\pi}'$ *t*-bisimulates $\vec{\pi} + \pi_d$ if, for all T with $|T| \leq t$ and inputs $\vec{x}$:*

- *for all adversaries $A = (\vec{\tau}_T, \sigma_e)$, there exists an adversary $A' = (\vec{\tau}'_T, \sigma'_e)$ such that $\vec{\pi}(\vec{x}, A)$ and $\vec{\pi}'(\vec{x}, A')$ are identically distributed;*
- *for all adversaries $A' = (\vec{\tau}'_T, \sigma'_e)$, there exists an adversary $A = (\vec{\tau}_T, \sigma_e)$ such that $\vec{\pi}(\vec{x}, A)$ and $\vec{\pi}'(\vec{x}, A')$ are identically distributed.*

Note that 0-bisimulation is equivalent to implementation. While implementation considers only what happens when there is no malicious behaviour, *t*-bisimulation generalizes this notion by taking malicious behavior into account. For some of our results (specifically, Theorems 4.1 and 4.2), we use this notion, and show that it is achievable under the conditions of these theorems.

We have required schedulers to eventually deliver each message that is sent. Because we assume that protocols in the mediator game are bounded, all protocols in the mediator game must terminate. This means that we can't hope to simulate a protocol in the cheap-talk game that deadlocks. (We assume that if the protocol deadlocks, it has a special output that we denote $\perp$. Given our constraints, we can never get an output of $\perp$ in the mediator game.) To deal with this situation, we relax this requirement on schedulers somewhat, but only in the mediator game. We take a *relaxed*

scheduler to be one that may never deliver some messages. There is no requirement on messages sent by the players. We can define *relaxed t -bisimulation* just as we defined t -bisimulation, except that we allow the schedulers σ'_e and σ_e to be relaxed schedulers. Finally, we define (t, t') -*bisimulation* just as we defined relaxed t -bisimulation except that σ'_e and σ_e must be standard if $|T| \leq t'$ (note that t' must be smaller than t).

We need a further property to deal with protocols that involve punishment strategies. For a punishment strategy to be effective, all the honest players have to play it. In our protocols, the punishment strategy is played when there is a deadlock (so some players never terminate); that is, the punishment strategy is in the honest players' "wills". Thus, we want it to be the case that either none of the honest players terminate (in which case the punishment strategy will be effective) or all of them do; we do not want it to be the case that only some of the honest players terminate.

Definition 5.2 ((t, t') -cotermination). *A protocol $\vec{\pi}$ (t, t') -coterminates if, for all schedulers σ_e , all subsets T of at most t players, and all strategy profiles $\vec{\tau}_T$ for the players in T , in all histories of the protocol $(\vec{\pi}_{-T}, \vec{\tau}_T, \sigma_e)$, either all the players not in T terminate, or at most t' players not in T do. We say that $\vec{\pi}$ t -coterminates if it $(t, 0)$ -coterminates.*

For some of our results, we need "approximate" versions of t -bisimulation, relaxed t -bisimulation, (t, t') bisimulation, t -cotermination and (t, t') -cotermination, that allow an ϵ probability of error. For t -bisimulation, relaxed t -bisimulation, and (t, t') bisimulation, this means that the distance between the distribution over outputs in the cheap-talk game and the distribution over outputs in the mediator game is at most ϵ (where the notion of distance is that used in the definition of ϵ -implementation in Section 2), while for t -cotermination and (t, t') -cotermination it means that the property holds with probability $1 - \epsilon$. We call these properties ϵ - t -bisimulation, relaxed ϵ - t -bisimulation, ϵ -(t, t')-bisimulation, ϵ - t -cotermination, and ϵ -(t, t')-cotermination respectively. In our companion paper [Geffner and Halpern 2018], we prove the following results:

Theorem 5.3. *Given a mediator game Γ_d extending Γ and a strategy profile $\vec{\sigma} + \sigma_d$ in Γ_d in canonical form, there exists a strategy profile $\vec{\sigma}_{CT}$ for Γ_{CT} such that $\vec{\sigma}_{CT}$ $(t, 2t + 1)$ -coterminates and t -bisimulates (resp., (t, t') -bisimulates) $\vec{\sigma} + \sigma_d$ if $t < n/3$ and $t < n/4$ (resp., $3t + t' < n$) respectively. If σ_d is responsive, then the number of messages sent in a history of $\vec{\sigma}_{CT}$ is polynomial in n and N , linear in c , and independent of $\vec{u}'$.*

Theorem 5.4. *Given a mediator game Γ_d extending Γ , a strategy profile $\vec{\sigma} + \sigma_d$ in Γ_d in canonical form, and a real number $\epsilon \in (0, 1]$, there exists a strategy profile $\vec{\sigma}_{CT}$ in Γ_{CT} such that $\vec{\sigma}_{CT}$ ϵ -($t, t + 1$)-coterminates and ϵ - t -bisimulates (resp., ϵ -(t, t')-bisimulates) $\vec{\sigma} + \sigma_d$ if $t < n/2$ and $t < n/3$ (resp., $2t + t' < n$) respectively. If σ_d is responsive, then the number of messages sent in a history of $\vec{\sigma}_{CT}$ is polynomial in n and N , linear in c , and independent of $\vec{u}'$.*

We use the constructions provided by Theorems 5.3 and 5.4 to prove Theorems 4.1, 4.2, 4.4, and 4.5. For Theorems 4.1 and 4.2, we show that these constructions already satisfy all the desired properties. For Theorems 4.4 and 4.5, $(t + k, t)$ -bisimulation and ϵ -($t + k, t$)-bisimulation guarantee that the only way in which rational and malicious players can make the outcome of the cheap-talk game different from that of the mediator game is by preventing the honest players from terminating. However, it is not in the interest of rational players that too many honest players end in deadlock, since if enough honest players do not terminate, they play the punishment strategy according to their wills.

We then show that the construction for Theorems 5.3 and 5.4 can be easily modified so that $\vec{\sigma}_{CT}$ terminate or none does. It would then seem that to always be in the interest of rational players that all honest players terminate. Unfortunately, a punishment strategy payoff is worse for rational players than the equilibrium payoff only in expectation, and thus, depending on the information rational players receive throughout the game, there might be some situations in which they might actually want to be punished (we provide a concrete example in Section 6.5). We deal with this issue by showing that all mediator games can be reduced to a mediator game where the mediator uses a *minimally-informative strategy*, so that players receive no information about the outcome until the very end. This makes the punishment strategy a persistent threat for rational players.

6 Proofs of the main theorems

6.1 Coordination between the environment and malicious players

Before proving the main results, it is useful to understand some of the implication of (k, t) -robustness, particularly when it comes to the interactions between the environment and the malicious and rational players. The definition of (k, t) -robustness requires that rational players have no profitable deviation, no matter what the malicious players and the environment do. It may seem *a priori* that the malicious players, the rational players, and the environment all act independently, but in fact, we can assume without loss of generality that they are all under the control of a single adversary. Clearly rational players can coordinate by sending messages to each other. The malicious and rational players can also coordinate with the environment so that, for example, the malicious and rational players can act knowing who will be scheduled when and the environment can schedule rational and malicious players based on their inputs. This follows from the fact that (k, t) -robustness must hold for all schedulers.

To see that a player i can communicate with the environment, recall that we have assumed that the message space is finite, say $\{m_0, \dots, m_M\}$. Immediately after sending m_j , i sends j empty messages to itself. So, even though the environment cannot read the messages, it will know that i sent message m_j .² (Clearly the environment will also know who sent the message, since the environment delivered the message.) Thus, we can assume without loss of generality that the rational and malicious players know the environment's protocol (and thus know when a message that is sent will be delivered), hence it suffices for the environment to tell the non-honest players when the k th message is sent, who sent it, and who the intended recipient is. All the non-honest players i initially send themselves $(n+1)^2$ empty messages. If the first message was sent by player j_1 to player j_2 (treating the mediator as player 0 in the mediator game), then the environment delivers $(n+1)j_1 + j_2$ of these empty messages. Then player i sends another $(n+1)^2$ empty messages, allowing the environment to encode the sender and receiver of the next message, if there is one.

The fact that the environment and the malicious players can communicate allows us to prove a tighter correspondence between deviations in the cheap-talk game and deviations in the mediator game than the one given by t -bisimulation.

²We can encode m_j using using fewer messages by having i send message to other agents as well as itself. Our goal here is not to minimize the number of messages, but to show that communication with the environment is possible.

Proposition 6.1. *Given two protocols $\vec{\pi}$ and $\vec{\pi}'$ and a scheduler σ_e , if $\vec{\pi}'$ t -bisimulates $\vec{\pi}$, there exists a function H_{σ_e} from strategies to strategies such that $H_{\sigma_e}(\vec{\pi}_i) = \vec{\pi}'_i$ for all i , and for all adversaries $A = (\vec{\tau}_T, \sigma_e)$ with $|T| \leq t$, there exists a scheduler σ'_e such that for all inputs $\vec{x}$, $\vec{\pi}(\vec{x}, A)$ and $\vec{\pi}'(\vec{x}, (H_{\sigma_e}(\vec{\tau}_T), \sigma'_e))$ are identically distributed (where we extend H_{σ_e} to strategy profiles by taking $H_{\sigma_e}(\tau_1, \dots, \tau_m) = (H_{\sigma_e}(\tau_1), \dots, H_{\sigma_e}(\tau_m))$).*

Proof. Since $\vec{\pi}'$ t -bisimulates $\vec{\pi} + \pi_d$, for each adversary $A = (\vec{\tau}_T, \sigma_e)$, there exists an adversary $A' = (\vec{\tau}'_T, \sigma'_e)$ such that $\vec{\pi}(\vec{x}, A)$ and $\vec{\pi}'(\vec{x}, A')$ are identically distributed for all inputs $\vec{x}$. This means that, fixing σ_e , there exists a well-defined function $H_{\sigma_e}^{adv}$ from strategy profiles to adversaries such that for all subsets T with $|T| \leq t$ and all strategies $\vec{\tau}_T$ for players in T , there exists a scheduler σ'_e such that $(\vec{\pi} + \pi_d)(\vec{x}, (\vec{\tau}_T, \sigma_e))$ and $\vec{\pi}'(\vec{x}, H_{\sigma_e}^{adv}(\vec{\tau}_T))$ for all inputs $\vec{x}$. Given an adversary $A = (\vec{\tau}_T, \sigma_e)$, consider the following adversary $A' = (\vec{\tau}'_T, \sigma'_e)$:

1. The scheduler begins the game by scheduling all players in T once.
2. Each player $i \in T$ sends the description of the strategy $\vec{\tau}_i$ to the scheduler σ'_e the first time it is scheduled.
3. After receiving the strategies of the players in T , the scheduler computes $H_{\sigma_e}^{adv}(\vec{\tau}_T)$ and sends each player $i \in T$ the (description of) strategy $H_{\sigma_e}^{adv}(\vec{\tau}_T)_i$, and then switches to using $H_{\sigma_e}^{adv}(\vec{\tau}_T)_e$.
4. Each player $i \in T$ switches to the strategy sent by the scheduler after receiving it.

Consider the function H_{σ_e} that maps τ_i to τ'_i if $\tau_i \neq \pi_i$ and maps π_i to π'_i (note that H_{σ_e} is well defined, since τ'_i does not depend on the strategy of other players in T) and a scheduler σ''_e that acts like σ'_e except that at step 1, it schedules only the players i in T such that $\tau_i \neq \pi_i$. We have by construction that $\vec{\pi}(\vec{x}, A) = \vec{\pi}'(\vec{x}, H_{\sigma_e}^{adv}(\vec{\tau}_T))$ and therefore that $\vec{\pi}(\vec{x}, A) = \vec{\pi}'(\vec{x}, (H_{\sigma_e}(\vec{\tau}_T), \sigma''_e))$, as desired. $\square$

Proposition 6.1 implies that we can assume without loss of generality that individual deviations in the mediator game correspond to individual deviations in the cheap-talk game and vice-versa. An analogous result holds for ϵ - t -bisimulation:

Proposition 6.2. *Given two protocols $\vec{\pi}$ and $\vec{\pi}'$ and a scheduler σ_e , if $\vec{\pi}'$ ϵ - t -bisimulates $\vec{\pi}$, there exists a function H_{σ_e} from strategies to strategies such that $H_{\sigma_e}(\vec{\pi}_i) = \vec{\pi}'_i$ for all i , and for all adversaries $A = (\vec{\tau}_T, \sigma_e)$ with $|T| \leq t$, there exists a scheduler σ'_e such that for all inputs $\vec{x}$, the distance between the distributions $\vec{\pi}(\vec{x}, A)$ and $\vec{\pi}'(\vec{x}, (H_{\sigma_e}(\vec{\tau}_T), \sigma'_e))$ is at most ϵ (where the notion of distance is that used in the definition of ϵ -implementation in Section 2).*

As we show next, since the scheduler can collude with malicious players, t -immune strategy profiles satisfy an even stronger condition: deviations by players in a set T with $|T| \leq t$ do not make things worse for the non-deviating players even if the environment colludes with the players in T .

Proposition 6.3. *If $\vec{\sigma}$ is t -immune, then for all sets T of players with $|T| \leq t$, strategies σ_e and σ'_e for the environment, strategy profiles $\vec{\tau}_T$ for the players in T , input profiles $\vec{x}$ and $\vec{x}'$, and players $i \notin T$, we have*

$$u_i(\Gamma_d, (\vec{\sigma}_{-T}, \vec{\tau}_T), \sigma'_e, \vec{x}'_T) \geq u_i(\Gamma_d, \vec{\sigma}, \sigma_e, \vec{x}_T). \quad (1)$$

Proof. Clearly, if (1) holds for all $\sigma_e, \sigma'_e, \vec{x}$, and $\vec{x}'$, then $\vec{\sigma}$ is t -immune. For the converse, suppose by way of contradiction that $\vec{\sigma}$ is t -immune but for some $T, \vec{\tau}, \sigma_e, \sigma'_e$, and $i \notin T$, (1) does not hold. Consider a scheduler σ''_e that acts just like σ_e , except that if some player i sends a message to itself it acts like σ'_e . Then players in T can effectively decrease i 's payoff with scheduler σ''_e by sending a message to themselves and playing as if they had input $\vec{x}'_T$; that is, there is a strategy $\vec{\tau}'_T$ such that

$$\begin{aligned} & u_i(\Gamma_d, (\vec{\sigma}_{-T}, \vec{\tau}'_T), \sigma''_e, \vec{x}_T) \\ &= u_i(\Gamma_d, (\vec{\sigma}_{-T}, \vec{\tau}_T), \sigma'_e, \vec{x}'_T) \\ &< u_i(\Gamma_d, \vec{\sigma}, \sigma_e, \vec{x}_T) \\ &= u_i(\Gamma_d, \vec{\sigma}, \sigma''_e, \vec{x}_T), \end{aligned}$$

contradicting the assumption that $\vec{\sigma}$ is t -immune. $\square$

A similar argument shows that (k, t) -robust strategy profiles satisfy a correspondingly stronger condition, made precise in the following proposition:

Proposition 6.4. *A strategy profile $\vec{\sigma}$ is (k, t) -robust (resp., strongly (k, t) -robust) if and only if it is t -immune and for all disjoint sets K and T with $1 \leq |K| \leq k$ and $|T| \leq t$, all strategy profiles $\vec{\tau}_K, \vec{\tau}_T$, and $\vec{\tau}'_T$ for the players in K and T , respectively, all environment strategies σ_e and σ'_e , and all input profiles $\vec{x}$ and $\vec{x}'$, we have that*

$$\begin{aligned} & u_i(\Gamma_d, (\vec{\sigma}_{-(K \cup T)}, \vec{\tau}_K, \vec{\tau}'_T), \sigma'_e, \vec{x}'_{(K \cup T)}) \\ & \leq u_i(\Gamma_d, (\vec{\sigma}_{-(K \cup T)}, \vec{\tau}_K, \vec{\tau}_T), \sigma'_e, \vec{x}'_{(K \cup T)}) \end{aligned} \quad (2)$$

for some $i \in K$ (resp., for all $i \in K$).

Proof. Again, it is clear that if (2) holds for all K and T with $1 \leq |K| \leq k$ and $|T| \leq t$, all $\vec{\tau}_K, \vec{\tau}_T, \vec{\tau}'_T, \vec{x}$, and $\vec{x}'$, and some (resp., all) $i \in K$, then $\vec{\sigma}$ is (k, t) -robust (resp., strongly (k, t) -robust).

For the converse, assume by way of contradiction that $\vec{\sigma}$ is (k, t) -robust, but for some disjoint sets K and T with $1 \leq |K| \leq k$ and $|T| \leq t$, $\vec{\tau}_K, \vec{\tau}_T, \vec{\tau}'_T, \vec{x}$, and $\vec{x}'$, and all $i \in K$, (2) does not hold. Again, we use the fact that the rational players can effectively communicate with malicious players and with the scheduler. Consider a scheduler σ''_e that acts like σ_e unless some player sends a message to itself, in which case it acts like σ'_e , and a strategy profile $\vec{\tau}''_T$ in which each player $i \in T$ acts as if it was using strategy $(\tau_T)_i$, except that it switches to $(\tau'_T)_i$ and acts as if it has input x'_i if it receives a message from a rational player (i.e., a player in K) asking it to do so. Then, given input profile $\vec{x}$, strategy profile $\vec{\tau}''_T$ for T , and scheduler σ''_e , player i can gain by sending a message to itself and sending a message to players in T asking them to follow $\vec{\tau}''_T$ and to act as if they have input $\vec{x}'_T$, and by having players in K play $\vec{\tau}_K$ as if they had input $\vec{x}'_K$, rather than playing $\vec{\sigma}$. This contradicts the assumption that $\vec{\sigma}$ is (k, t) -robust. The argument for strong (k, t) -robustness is analogous. $\square$

Another property interesting in its own right that follows from this argument is that (k, t) -robust strategies must be *scheduler-proof*: the expected payoff for all players is the same regardless of the scheduler:

Corollary 6.5. *If $\vec{\sigma}$ is (k, t) -robust for some $k \geq 1$, then for all sets T with $|T| \leq t$, strategy profiles $\vec{\tau}_T$ for the players in T , environment strategies σ_e and σ'_e , input profiles $\vec{x}$, and players $i \notin T$, we have $u_i(\Gamma, (\vec{\sigma}_{-T}, \vec{\tau}_T), \sigma_e, \vec{x}_T) = u_i(\Gamma, (\vec{\sigma}_{-T}, \vec{\tau}_T), \sigma'_e, \vec{x}_T)$.*

We have analogous strengthenings of ϵ - t -immunity and ϵ -(k, t)-robustness, which are stated next. The proofs are essentially identical to that of Proposition 6.3 and 6.4 respectively, so we omit them here.

Proposition 6.6. *If $\epsilon > 0$ and $\vec{\sigma}$ is ϵ - t -immune in game Γ , then for all sets T of players with $|T| \leq t$, strategy profiles $\vec{\tau}_T$ for the players in T , environment strategies σ_e and σ'_e , input profiles $\vec{x}$ and $\vec{x}'$, and players $i \notin T$, we have that*

$$u_i(\Gamma, (\vec{\sigma}_{-T}, \vec{\tau}_T), \sigma'_e, \vec{x}') > u_i(\Gamma, \vec{\sigma}, \sigma_e, \vec{x}) - \epsilon.$$

Proposition 6.7. *A strategy profile $\vec{\sigma}$ is ϵ -(k, t)-robust (resp., strongly ϵ -(k, t)-robust) in game Γ if and only if it is ϵ - t -immune and, for all disjoint sets K and T of players with $1 \leq |K| \leq k$ and $|T| \leq t$, all strategy profiles $\vec{\tau}_K$, $\vec{\tau}'_T$, and $\vec{\tau}_T$ for players in K and T , respectively, all environment strategies σ_e and σ'_e , and all input profiles $\vec{x}$ and $\vec{x}'$, we have that*

$$u_i(\Gamma, (\vec{\sigma}_{-(K \cup T)}, \vec{\tau}_K, \vec{\tau}'_T), \sigma'_e, \vec{x}'_{(K \cup T)}) < u_i(\Gamma, (\vec{\sigma}_{-T}, \vec{\tau}_T), \sigma_e, \vec{x}_T) + \epsilon$$

for some $i \in K$ (resp., for all $i \in K$).

It will be useful for our later results that we can actually improve on the bound of ϵ in Propositions 6.6 and 6.7.

Proposition 6.8. *If $\vec{\sigma}$ is an ϵ - t -immune strategy in a finite game Γ , then there exists ϵ_0 with $0 < \epsilon_0 < \epsilon$ such that for all sets of players T with $|T| \leq t$, strategy profiles $\vec{\tau}_T$ for the players in T , environment strategies σ_e and σ'_e , input profiles $\vec{x}$ and $\vec{x}'$, and players $i \notin T$, we have that*

$$u_i(\Gamma, (\vec{\sigma}_{-T}, \vec{\tau}_T), \sigma'_e, \vec{x}') > u_i(\Gamma, \vec{\sigma}, \sigma_e, \vec{x}) - \epsilon_0.$$

Proof. Since, by Proposition 6.6, for each choice of $\vec{\tau}_T$, σ_e , and σ'_e , we have

$$u_i(\Gamma, \vec{\sigma}, \sigma_e, \vec{x}) - u_i(\Gamma, (\vec{\sigma}_{-T}, \vec{\tau}_T), \sigma'_e, \vec{x}') < \epsilon,$$

and the space of player strategy profiles, environment strategies, and input value profiles is compact, if we take the sup of the left-hand side over all choices of strategy profiles $\vec{\tau}_T$, environment strategies σ_e and σ'_e , and input profiles $\vec{x}$ and $\vec{x}'$, it takes on some maximum value $\epsilon_1 < \epsilon_0$. We can then take $\epsilon_0 = (\epsilon + \epsilon_1)/2$. $\square$

Using Proposition 6.7, we get a similar result for ϵ -(k, t)-robustness. The proof is analogous to that of Proposition 6.8.

Proposition 6.9. *If Γ is a finite game and $\vec{\sigma}$ is a ϵ -(k, t)-robust strategy (resp., strongly ϵ -(k, t)-robust strategy) in Γ_d , then there exists ϵ_0 with $0 < \epsilon_0 < \epsilon$ such that for all disjoint sets K and T of players with $1 \leq |K| \leq k$ and $|T| \leq t$, all strategy profiles $\vec{\tau}_K$, $\vec{\tau}_T$ and $\vec{\tau}'_T$ for players in K and T , respectively, all environment strategies σ_e and σ'_e , and all input profiles $\vec{x}$ and $\vec{x}'$, we have that*

$$u_i(\Gamma, (\vec{\sigma}_{-(K \cup T)}, \vec{\tau}_K, \vec{\tau}_T), \sigma_e, \vec{x}'_{(K \cup T)}) < u_i(\Gamma, (\vec{\sigma}_{-T}, \vec{\tau}'_T), \sigma'_e, \vec{x}_T) + \epsilon_0$$

for some $i \in K$ (resp., all $i \in K$).

6.2 Constructing a protocol that t -coterminates

If $t < n/3$ we can get an analogue of Theorems 5.3 and 5.4 by replacing $(t, 2t+1)$ -cotermination by t -cotermination and ϵ -($t, t+1$)-cotermination by ϵ - t -cotermination respectively. We sketch the construction for t -cotermination; an analogous construction achieves ϵ - t -cotermination.

Consider a protocol $\vec{\sigma}'_{CT}$ in which players play just as in $\vec{\sigma}_{CT}$ except that, whenever an honest player i terminates in $\vec{\sigma}_{CT}$, it instead broadcasts an ‘OK’ message to all players and waits until it receives $3t+1$ ‘OK’ messages before it terminates in $\vec{\sigma}'_{CT}$. Note that if an honest player terminates in $\vec{\sigma}'_{CT}$, then at least $3t+1$ players must have broadcast an ‘OK’ message in this history of $\vec{\sigma}'_{CT}$, of which at least $2t+1$ are honest. Thus, at least $2t+1$ players terminate in the corresponding history of $\vec{\sigma}_{CT}$. Since $\vec{\sigma}_{CT}$ $(t, 2t+1)$ -coterminates, it follows that all players not in T must terminate with $\vec{\sigma}_{CT}$, and hence all players not in T send an ‘OK’ message (and terminate) with $\vec{\sigma}'_{CT}$.

It remains to show that $\vec{\sigma}'_{CT}$ still (t, t') -bisimulates $\vec{\sigma} + \sigma_d$ if $3t + t' < n$. Clearly, it still relaxed t -bisimulates $\vec{\sigma} + \sigma_d$, so we just have to show that all players are guaranteed to terminate in the presence of at most t' malicious players. However, in this case, by assumption, all honest players are guaranteed to terminate in $\vec{\sigma}_{CT}$, and thus all honest players are guaranteed to eventually send an ‘OK’ broadcast in $\vec{\sigma}'_{CT}$. Since $n - t' > 3t + 1$, this guarantees that there will be at least $3t + 1$ ‘OK’ broadcasts and all honest players will eventually terminate, as desired.

Note that this construction requires players a reliable broadcast protocol, and thus can be done only if $n > 3(t + k)$. To prove Theorem 4.5 we require different techniques.

6.3 Proof of Theorem 4.1

By Theorem 5.3, if $n > 4k + 4t$, there exists a strategy profile $\vec{\sigma}_{CT}$ that $(k + t)$ -bisimulates $\vec{\sigma} + \sigma_d$. It is immediate from the definition of $(k + t)$ -bisimulation that $\vec{\sigma}_{CT}$ implements $\vec{\sigma} + \sigma_d$. Since the probability of deadlock is 0, what the players do in case of deadlock is irrelevant, so this approach works both in the case of the AH approach and the default-move approach. It remains to show that, for each utility variant $\Gamma_d(\vec{u}')$ of Γ_d , if $\vec{\sigma} + \sigma_d$ is a (strongly) (k, t) -robust equilibrium in $\Gamma_d(\vec{u}')$, then $\vec{\sigma}_{CT}$ is a (strongly) (k, t) -robust equilibrium in $\Gamma_{CT}(\vec{u}')$. We start by showing that $\vec{\sigma}_{CT}$ is t -resilient in $\Gamma_{CT}(\vec{u}')$.

Given T with $|T| \leq t$, $\vec{\tau}_T$, and σ_e , by Theorem 5.3 and Proposition 6.1, there exists a function H_{σ_e} from strategies to strategies and a scheduler σ'_e such that for all input profiles $\vec{x}$,

$$\begin{aligned} & u'_i(\Gamma_{CT}(\vec{u}'), ((\vec{\sigma}_{CT})_{-T}, \tau_T), \sigma_e, \vec{x}) \\ &= u'_i(\Gamma_d(\vec{u}'), (\vec{\sigma}_{-T}, H_{\sigma_e}(\vec{\tau}_T)), \sigma'_e, \vec{x}) \end{aligned}$$

for all players i . There also exists a scheduler σ''_e such that

$$u'_i(\Gamma_{CT}(\vec{u}'), \vec{\sigma}_{CT}, \sigma'_e, \vec{x}) = u'_i(\Gamma_d(\vec{u}'), \vec{\sigma}, \sigma''_e, \vec{x}).$$

Since $\vec{\sigma}$ is t -immune, for all $i \notin T$ we have that

$$\begin{aligned} & u'_i(\Gamma_{CT}(\vec{u}'), ((\vec{\sigma}_{CT})_{-T}, \vec{\tau}_T), \sigma_e, \vec{x}_T) \\ &= u'_i(\Gamma_d(\vec{u}'), (\vec{\sigma}_{-T}, H_{\sigma_e}(\vec{\tau}_T)), \sigma'_e, \vec{x}_T) \\ &\geq u'_i(\Gamma_d(\vec{u}'), \vec{\sigma}, \sigma''_e, \vec{x}_T) \quad [\text{by Lemma 6.3}] \\ &= u'_i(\Gamma_{CT}(\vec{u}'), \vec{\sigma}_{CT}, \sigma'_e, \vec{x}_T). \end{aligned}$$

Therefore, $\vec{\sigma}_{CT}$ is t -immune.

To show (strong) (k, t) -robustness, taking $\vec{\tau}_T$, σ_e , and σ'_e as above, suppose that K is a set of players disjoint from T such that $|K| \leq k$, and the players in K play $\vec{\tau}_K$. By Theorem 5.3 and Proposition 6.1, there exists σ_e^* and H_{σ_e} such that

$$\begin{aligned} & u'_i(\Gamma_{CT}(\vec{u}'), ((\vec{\sigma}_{CT})_{-(K \cup T)}, \vec{\tau}_K, \vec{\tau}_T), \sigma_e, \vec{x}) \\ = & u'_i(\Gamma_d(\vec{u}'), (\vec{\sigma}_{-(K \cup T)}, H_{\sigma_e}(\vec{\tau}_K), H_{\sigma_e}(\vec{\tau}_T)), \sigma_e^*, \vec{x}_{(K \cup T)}) \end{aligned}$$

for all players i . By Corollary 6.4, if $\vec{\sigma} + \sigma_d$ is (k, t) -robust (resp., strongly (k, t) -robust) in $\Gamma_d(\vec{u}')$, then

$$\begin{aligned} & u'_i(\Gamma_d(\vec{u}'), (\vec{\sigma}_{-(K \cup T)}, H_{\sigma_e}(\vec{\tau}_K), H_{\sigma_e}(\vec{\tau}_T)), \sigma_e^*, \vec{x}_{(K \cup T)}) \\ \leq & u'_i(\Gamma_d(\vec{u}'), (\vec{\sigma}_{-T}, H_{\sigma_e}(\vec{\tau}_T)), \sigma'_e, \vec{x}_T) \end{aligned}$$

for some (resp., all) $i \in K$. For those $i \in K$ for which this inequality holds, we have

$$\begin{aligned} & u'_i(\Gamma_{CT}(\vec{u}'), ((\vec{\sigma}_{CT})_{-(K \cup T)}, \tau_K, \tau_T), \sigma_e, \vec{x}_{(K \cup T)}) \\ = & u'_i(\Gamma_d(\vec{u}'), (\vec{\sigma}_{-(K \cup T)}, H_{\sigma_e}(\vec{\tau}_K), H_{\sigma_e}(\vec{\tau}_T)), \sigma_e^*, \vec{x}_{(K \cup T)}) \\ \leq & u'_i(\Gamma_d(\vec{u}'), (\vec{\sigma}_{-T}, H_{\sigma_e}(\vec{\tau}_T)), \sigma'_e, \vec{x}_T) \\ = & u'_i(\Gamma_{CT}(\vec{u}'), ((\vec{\sigma}_{CT})_{-T}, \sigma_e, \vec{x}_T)). \end{aligned}$$

It follows that $\vec{\sigma}_{CT}$ is (strongly) (k, t) -robust in $\Gamma_{CT}(\vec{u}')$.

6.4 Proof of Theorem 4.2

The proof of Theorem 4.2 is essentially the same as that of Theorem 4.1, except that we now use Theorem 5.4 instead of Theorem 5.3. By Theorem 5.4, for all $\epsilon' \in (0, 1]$, there exists a protocol $\vec{\sigma}_{CT}$ that ϵ -($t + k$)-bisimulates $\vec{\sigma}$ and the expected number of messages sent is polynomial in n and N , and linear in c . It follows that $\vec{\sigma}_{CT}$ ϵ' -implements $\vec{\sigma}$ and has at most a probability ϵ' of deadlock. We next show that the can make ϵ' sufficiently small so that the question of whether we use AH approach or the default-move approach becomes irrelevant.

We now prove ϵ -(k, t)-robustness. Suppose that $\vec{\sigma} + \sigma_d$ is a (strongly) ϵ -(k, t)-robust equilibrium in the utility variant $\Gamma_d(\vec{u}')$ of Γ_d . We show that $\vec{\sigma}_{CT}$ is ϵ - t -immune in $\Gamma_{CT}(\vec{u}')$. Since $\vec{\sigma}_{CT}$ ϵ -($t + k$)-bisimulates $\vec{\sigma}$, for all inputs $\vec{x}$ we can associate histories in the mediator game and histories in the cheap-talk game in such a way that the set of histories where the outcomes differ has probability at most ϵ' . Since all utilities are in the range $[-M/2, M/2]$, by assumption, the maximum difference in utility between two outcomes in the underlying game is M . Thus, by Proposition 6.2, there exists an environment strategy σ'_e and a function H_{σ_e} from strategies to strategies such that for all input profiles $\vec{x}$, we have

$$\begin{aligned} & u'_i(\Gamma_{CT}(\vec{u}'), ((\vec{\sigma}_{CT})_{-T}, \vec{\tau}_T), \sigma'_e, \vec{x}) \\ > & u'_i(\Gamma_d(\vec{u}'), (\vec{\sigma}_{-T}, H_{\sigma_e}(\vec{\tau}_T)), \sigma_e, \vec{x}_T) - \epsilon' M \end{aligned}$$

for all $i \notin T$. Theorem 5.4 also guarantees that there exists an environment strategy σ''_e such that

$$u'_i(\Gamma_{CT}(\vec{u}'), \vec{\sigma}_{CT}, \sigma'_e, \vec{x}_T) < u'_i(\Gamma_d(\vec{u}'), \vec{\sigma}, \sigma''_e, \vec{x}_T) + \epsilon' M.$$

Since $\vec{\sigma}$ is ϵ - t immune in $\Gamma_d(\vec{u}')$, by Proposition 6.8, there exists a value ϵ_0 with $0 < \epsilon_0 < \epsilon$ such that

$$\begin{aligned} & u'_i(\Gamma_{CT}(\vec{u}'), ((\vec{\sigma}_{CT})_{-T}, \vec{\tau}'_T), \sigma'_e, \vec{x}_T) \\ & > u'_i(\Gamma_d(\vec{u}'), (\vec{\sigma}_{-T}, H_{\sigma_e}(\vec{\tau}'_T)), \sigma_e, \vec{x}_T) - \epsilon' M \\ & > u'_i(\Gamma_d(\vec{u}'), \vec{\sigma}, \sigma'_e, \vec{x}_T) - \epsilon_0 - \epsilon' M \\ & > u'_i(\Gamma_{CT}(\vec{u}'), \vec{\sigma}_{CT}, \sigma'_e, \vec{x}_T) - \epsilon_0 - 2\epsilon' M. \end{aligned}$$

If we take $\epsilon' = (\epsilon - \epsilon_0)/2M$, this shows that $\vec{\sigma}_{CT}$ is (ϵ, t) -immune with both the AH approach and the default-move approach.

To show (strong) ϵ -(k, t)-robustness, keeping T , $\vec{\tau}_T$, H_{σ_e} , σ_e , and σ'_e as above, for all sets K of players disjoint from T with $1 \leq |K| < k$ and strategy profiles $\vec{\tau}_K$, there exists an environment strategy σ_e^* and a value ϵ_0 with $0 < \epsilon_0 < \epsilon$ such that for all input profiles $\vec{x}$, if $\vec{\sigma} + \sigma_d$ is (k, t) -robust (resp. strongly (k, t) -robust), then

$$\begin{aligned} & u'_i(\Gamma_{CT}(\vec{u}'), ((\vec{\sigma}_{CT})_{-(K \cup T)}, \vec{\tau}_K, \vec{\tau}_T), \sigma'_e, \vec{x}_{(K \cup T)}) \\ & < u'_i(\Gamma_d(\vec{u}'), (\vec{\sigma}_{-(K \cup T)}, H_{\sigma_e}(\vec{\tau}_K), H_{\sigma_e}(\vec{\tau}_T)), \sigma_e, \vec{x}_{(K \cup T)}) + \epsilon' M \\ & < u'_i(\Gamma_d(\vec{u}'), (\vec{\sigma}_{-T}, H_{\sigma_e}(\vec{\tau}_T)), \sigma'_e, \vec{x}_T) + \epsilon_0 + \epsilon' M & \text{[by Proposition 6.9]} \\ & < u'_i(\Gamma_{CT}(\vec{u}'), (\vec{\sigma}_{CT})_{-T}, \vec{\tau}, \sigma_e, \vec{x}_T) + \epsilon_0 + 2\epsilon' M & \text{[by Theorem 5.4].} \end{aligned}$$

for some (resp., for all) $i \in K$. This shows that if we take $\epsilon' := (\epsilon - \epsilon_0)/2M$, then $\vec{\sigma}_{CT}$ is ϵ -(k, t)-robust (resp., strongly (k, t) -robust). Note that this argument works for both the AH approach and the default-move approach since it does not depend on the actions played in deadlock.

6.5 Proof of Theorem 4.4

The proof of Theorem 4.4 is similar in spirit to that of Theorem 4.1. The main problem we have to deal with is that of ensuring that rational players participate. To force participation, we have the honest players put the punishment strategy in their “wills”, so that if $\vec{\sigma}_{CT}$ ends in deadlock, the rational players will be punished. By the arguments given in Section 6.2, we can assume without loss of generality in this proof that the implementation given by Theorems 5.4 ($t+k$)-coterminates, and thus either all honest players terminate or they all play the punishment strategy. Unfortunately, a naive implementation of this approach does not work, as the following example shows.

Consider an underlying game Γ for $n > 3k$ players where the set of actions is $A := \{0, 1, \perp\}$. If at least $k+1$ players play $\perp$, all players get a payoff of 1.1; if k or fewer players play $\perp$ and all players play either 0 or $\perp$, then all players get a payoff of 1; if k or fewer players play $\perp$ and all players play either $\perp$ or 1, then all players get a payoff of 2; otherwise, all players get 0. Let Γ_d be an extension of Γ with a mediator. Suppose that the mediator d uses the following strategy: The mediator d chooses $a, b \in \{0, 1\}$ uniformly at random. Then d sends the message $a + bi \pmod{2}$ to player i (the same a and b are used in all these messages). Finally, d sends the message “output b ; STOP” to all players (so the strategy is in canonical form).

Let σ_i be the strategy where player i ignores the message $a + bi$ and plays b after receiving the message “output b ”. It is easy to check that $\vec{\sigma}$ is a k -resilient equilibrium in the mediator game, and gives players an expected payoff of 1.5. Moreover, playing $\perp$ is a k -punishment strategy with respect to $\vec{\sigma}$, since if all but k players play $\perp$, then everyone gets a payoff of 1.1 (since at least $k+1$ players play $\perp$), which is less than 1.5.

The naive approach to implementing the mediator does not work for this game, at least with the punishment strategy $\perp$. For example, suppose that after receiving the messages $a + bi \pmod{2}$, the rational players communicate with each other. Moreover, suppose that the set K of rational players includes i and j such that $i - j$ is odd. Then the rational players can compute b . If $b = 0$, they actually prefer their payoff with the punishment strategy to their payoff with $\vec{\sigma}_{CT}$. Thus, they will stop sending messages. The simulation will not terminate, so the punishment strategy in the players' wills will be applied, making the rational players better off. Thus, we cannot simulate the mediator with this approach. Of course, there are punishment strategies in this game that would lead to cooperation (e.g., randomizing between 0 and 1). Nevertheless, this example shows that using an arbitrary punishment strategy may not suffice to force the rational players to cooperate.

The problem here is that the mediator tells each player i what $a + bi$ is. We do not want the mediator to send such unnecessary information. But what counts as unnecessary? We make “unnecessary” precise by showing that, for each strategy profile $\vec{\sigma} + \sigma_d$ of a mediator game, we can construct a strategy $\vec{\sigma}^m + \sigma_d^m$ that implements $\vec{\sigma} + \sigma_d$ and leaks no information. More precisely, there exists a function f from strategy profiles to strategy profiles such that, for all strategy profiles $\vec{\sigma} + \sigma_d$, $f(\vec{\sigma} + \sigma_d)$ implements $\vec{\sigma} + \sigma_d$ and essentially all the mediator sends each player when playing $f(\vec{\sigma} + \sigma_d)$ is the action to play in the underlying game. (If we require only weak implementation, then this is exactly the case; for implementation, the messages can also include a round number.) Moreover, if $\vec{\sigma} + \sigma_d$ is (k, t) -robust (resp., strongly (k, t) -robust, ϵ -(k, t)-robust, strongly ϵ -(k, t)-robust), then so is $f(\vec{\sigma} + \sigma_d)$. The construction of f proceeds as follows:

Let $D(\vec{x}, \sigma_e)$ be the distribution over action profiles in the underlying game that when playing $\vec{\sigma} + \sigma_d$ with input $\vec{x}$ and scheduler σ_e . The intuition behind the construction of $\vec{\sigma}^m + \sigma_d^m := f(\vec{\sigma} + \sigma_d)$ is that all players send their input to the mediator, the mediator waits until it receives messages from at least $n - k - t$ players, then simulates the game using the inputs sent by the players, and sends back to each player what they would have played in the simulation. However, to get an implementation of $\vec{\sigma} + \sigma_d$, the scheduler that the mediator uses in its simulation must depend somehow on the actual scheduler and must be chosen in such a way that, for a given input profile $\vec{x}$, all distributions in $\{D(\vec{x}, \sigma_e)\}_{\sigma_e}$ are possible.

More precisely, the construction proceeds as follows: player i uses strategy σ_i^m , according to which i sends input x_i to the mediator, waits for the mediator's message msg_i (which we take to be an action for player i), and then plays action msg_i . The mediator's strategy σ_d^m consists of waiting until the first turn ℓ at which there exists a subset S of at least $n - k - t$ players such that the mediator has received exactly one message s_i from each player $i \in S$, and s_i is a possible input of player i . What the mediator does next depends on whether $|S| = n$ or $|S| < n$.

If $|S| < n$, the mediator simulates $\vec{\sigma} + \sigma_d$ assuming that each player i in S has input s_i , and that the scheduler schedules players in S and the mediator in round-robin fashion and delivers all messages immediately after they are sent (note that such a scheduler exists even with the constraint that all players must be eventually scheduled, since the scheduler can schedule players not in S after the mediator terminates). The mediator then sends to each player i the action that i plays in its simulation. For future reference, we denote the scheduler used in this simulation by σ_e^S .

If $|S| = n$, the mediator proceeds as follows: Let Ω be the set of deterministic schedulers, and let $\mathcal{D}^{\vec{x}} := \bigcup_{\sigma_e \in \Omega} \{D(\vec{x}, \sigma_e)\}$. Since each player uses a finite amount of randomization in $\vec{\sigma} + \sigma_d$, for all inputs $\vec{x}$, deterministic schedulers σ_e , and action profiles $\vec{a}$, the probability that players play action profile $\vec{a}$ in the underlying game when playing $\vec{\sigma} + \sigma_d$ with input $\vec{x}$ and scheduler

σ_e is a rational number. Since there are finitely many possible action profiles in the underlying game, $\mathcal{D}^{\vec{x}}$ is countable. Thus, there exists a set $\{\sigma_e^{(\vec{x},1)}, \sigma_e^{(\vec{x},2)}, \dots\}$ of schedulers such that $\mathcal{D}^{\vec{x}} = \{D(\vec{x}, \sigma_e^{(\vec{x},1)}), D(\vec{x}, \sigma_e^{(\vec{x},2)}), \dots\}$. The mediator simulates $\vec{\sigma} + \sigma_d$ assuming that each player i has input s_i and that the scheduler is $\sigma_e^{(\vec{s}, \ell)}$ (recall that ℓ is the turn at which the mediator receives the required number of messages). If $\mathcal{D}^{\vec{x}}$ is finite, it performs the simulation with scheduler $\sigma_e^{(\vec{s}, \ell(\bmod |\mathcal{D}^{\vec{x}}|))}$ instead.

Lemma 6.10. $\vec{\sigma}^m + \sigma_d^m$ implements $\vec{\sigma} + \sigma_d$ and is (k, t) -robust.

Proof. First note that it suffices to prove this result for deterministic schedulers. This follows from the fact that all randomized schedulers can be written as a (possibly infinite) linear combination of deterministic schedulers. By construction, for all inputs $\vec{x}$ and all deterministic schedulers $\vec{x}$ (under the assumption that no agents deviate). To prove the converse, given a deterministic scheduler σ_e for $\vec{\sigma} + \sigma_d$, we have that $D(\vec{x}, \sigma_e) \in \mathcal{D}^{\vec{x}}$, and thus there exists $k \in \mathbb{N}$ such that $D(\vec{x}, \sigma_e) = D(\vec{x}, \sigma_e^{(\vec{x}, k)})$. Consider a scheduler σ'_e in $\vec{\sigma}^m + \sigma_d^m$ that schedules all honest players consecutively, then schedules the mediator $k - 1$ times, then delivers all messages sent by the players to the mediator, and then schedules the mediator again. By construction, in this scenario, the mediator simulates $\vec{\sigma} + \sigma_d$ with input profile $\vec{x}$ and scheduler $\sigma_e^{(\vec{x}, k)}$. Therefore $(\vec{\sigma} + \sigma_d)(\vec{x}, \sigma_e)$ and $(\vec{\sigma}^m + \sigma_d^m)(\vec{x}, \sigma'_e)$ are identically distributed.

To see that $\vec{\sigma}^m + \sigma_d^m$ is t -immune, suppose, by way of contradiction, that it is not. Thus, there must exist an adversary $A = (T, \vec{\tau}_T, \sigma_e)$ with $|T| \leq t$ and an input profile $\vec{x}$ such that $u_i(\vec{\sigma}^m + \sigma_d^m, A, \vec{x}) < u_i(\vec{\sigma} + \sigma_d, \vec{x}, \sigma_e)$ for some $i \notin T$. We can assume without loss of generality that A is deterministic. When playing $\vec{\sigma}^m + \sigma_d^m$ with adversary A and input profile $\vec{x}_T$, the first round m by which the mediator receives messages of the right form from a subset S of at least $n - k - t$ players, the set S and the values s_i used in the mediator's simulation are uniquely determined. Consider an adversary $A' = (T, \vec{\tau}'_T, \sigma'_e)$ in $\vec{\sigma} + \sigma_d$ where each player $i \in T$ acts as if it was an honest player with input s_i . while the scheduler acts like σ_e^S if $|S| < n$, and like σ_e^m otherwise. However, if a_i is the action that i would have played if it was honest and had input s_i , instead of playing a_i , i plays what it would have played in $\vec{\sigma}^m + \sigma_d^m$ if it had received message a_i from the mediator. By construction, for all inputs $\vec{x}$, $(\vec{\sigma} + \sigma_d)(\vec{x}, A)$ and $(\vec{\sigma}^m + \sigma_d^m)(\vec{x}, A')$ are identically distributed. This implies that

$$u_i(\vec{\sigma} + \sigma_d, A', \vec{x}) < u_i(\vec{\sigma} + \sigma_d, \vec{x}, \sigma_e'')$$

for some scheduler σ_e'' , which contradicts the assumption that $\vec{\sigma} + \sigma_d$ is t -immune.

The argument that $\vec{\sigma}^m + \sigma_d^m$ is (k, t) -resilient is analogous, and left to the reader. $\square$

Thus, without loss of generality, we can assume that the players and mediator use such a strategy profile. We call $f(\vec{\sigma} + \sigma_d)$ the *minimally-informative* strategy corresponding to $\vec{\sigma} + \sigma_d$. More generally, we say that $\vec{\sigma}^m + \sigma_d^m$ is a minimally-informative strategy if $\vec{\sigma}^m + \sigma_d^m = f(\vec{\sigma} + \sigma_d)$ for some strategy profile $\vec{\sigma} + \sigma_d$.

Since we consider only mediator games in canonical form, this guarantees termination for all honest players regardless of what the rational and malicious players do, provided that the scheduler is standard (i.e., not relaxed). However, once we allow relaxed schedulers, there is a possibility of deadlock. We assume for the purposes of the proof that we use the AH approach in the mediator game, and have the players play the punishment strategy in their wills. Since $\vec{\sigma}_{CT}$

guarantees t -cotermination for $t < n/3$, it follows that in the cheap-talk game, either all honest players terminate or all honest players play the punishment strategy. This guarantees that the players get the same payoff in corresponding histories in the mediator game and the cheap-talk game.

The next step in proving Theorem 4.4 is to show that rational players playing with a relaxed scheduler cannot get an expected payoff that is higher than their expected payoff when they play such a minimally-informative (k, t) -robust equilibrium strategy with a standard scheduler. Although this property does not hold in general, it does hold when a certain degree of cotermination (which is provided by Theorems 5.3 and 5.4) is guaranteed and there exists a punishment strategy. Under these conditions, the rational players do not want too many honest players to fail to terminate, because the honest players that do not terminate will play the punishment strategy.

To state this more precisely, we need to introduce a little more notation. Bisimulation guarantees that for each strategy τ_A that the adversary can play in the cheap-talk game, there exists a corresponding strategy τ'_A in the mediator game that leads to the same outcome for all players, regardless of the input. Since the strategy $\vec{\sigma}_{CT}$ provided by Theorems 5.3 and 5.4 coterminates (with the parameters of cotermination depending on the theorem), this imposes a constraint on τ'_A that is captured in the following definition:

Definition 6.11. *Given a strategy profile $\vec{\sigma}$, a scheduler σ_e , and a subset T of players, $(\vec{\sigma}, \sigma_e)$ T - t -coterminates if, for all input profiles $\vec{x}$, in every history of $(\vec{\sigma}, \sigma_e, \vec{x})$, either all players not in T terminate or less than t do. We say that $(\vec{\sigma}, \sigma_e)$ ϵ - T - t -coterminates if this property holds with probability $1 - \epsilon$.*

Proposition 6.12. *If $\epsilon > 0$, $\vec{\sigma} + \sigma_d$ is a minimally-informative ϵ -(k, t)-robust (resp., strongly ϵ -(k, t)-robust) equilibrium in a mediator game Γ_d for which a $(2k + 2t)$ -punishment strategy exists, σ_E is a relaxed scheduler, K and T are disjoint sets of players with $1 \leq |K| \leq k$ and $|T| \leq t$, and $\vec{\tau}_{(K \cup T)}$ is a strategy profile for the players in $K \cup T$ such that $(\vec{\sigma}_{-(K \cup T)}, \vec{\tau}_{(K \cup T)}, \sigma_E)$ ϵ -($K \cup T$)-($t + k + 1$)-coterminates, then there exists a value $\epsilon_0 < \epsilon$ such that for all standard schedulers σ_e and all input profiles $\vec{x}$, we have that*

$$\begin{aligned} & u'_i(\Gamma_d, (\vec{\sigma}_{-(K \cup T)}, \vec{\tau}_{(K \cup T)}), \sigma_E, \vec{x}_{(K \cup T)}) \\ & < u'_i(\Gamma_d, (\vec{\sigma}_{-T}, \vec{\tau}_T), \sigma_e, \vec{x}_T) + \epsilon_0 + \epsilon M \end{aligned}$$

for some (resp., for all) $i \notin T$.

To prove Proposition 6.12, we first show that all strategy profiles can be approximated by a profile where there is a uniform bound on the amount of randomness used by an adversary.

Definition 6.13. *Given a strategy profile $\vec{\sigma}$, an adversary $A = (\vec{\tau}_T, \sigma_e)$ is N -bounded with respect to $\vec{\sigma}$ if, for all inputs and all histories, the number of random coin tosses performed by a player in T or the scheduler σ_e when players in T play $\vec{\tau}_T$, the remaining players play $\vec{\sigma}$, and the scheduler plays σ_e is bounded by N .*

Lemma 6.14. *For all strategy profiles $\vec{\sigma}$, adversaries $A = (\vec{\tau}_T, \sigma_E)$, and $\epsilon > 0$, there exists an adversary $A' = (\vec{\tau}'_T, \sigma'_E)$ and an $N > 0$ such that A' is N -bounded with respect to $\vec{\sigma}$ and, for all input profiles $\vec{x}$, the distance between the distributions $O(\vec{\sigma}, A, \vec{x})$ and $O(\vec{\sigma}, A', \vec{x})$ is at most ϵ .*

Proof. Fix $\epsilon > 0$. Let $A^N = (\vec{\tau}_T^N, \sigma_E^N)$ be the adversary that plays $(\vec{\tau}_T, \sigma_E)$, except that all players $i \in T$ and the scheduler act as if all the coin tosses after the N th coin toss are tails. By construction, for all input profiles $\vec{x}$, we have

$$\lim_{N \rightarrow \infty} d(O(\vec{\sigma}, A, \vec{x}), O(\vec{\sigma}, A^N, \vec{x})) = 0,$$

where $d(\cdot, \cdot)$ denotes the distance between distributions. Thus, there exists an integer $N_{\vec{x}}$ such that $d(O(\vec{\sigma}, A, \vec{x}), O(\vec{\sigma}, A^{N_{\vec{x}}}, \vec{x})) < \epsilon$. Since there are only finitely many input profiles, we can take $N = \max_{\vec{x}}(N_{\vec{x}})$ to get the desired result. $\square$

Proof of Proposition 6.12. First assume that

$A := (\vec{\tau}_{(K \cup T)}, \sigma_E)$ is N -bounded with respect to $\vec{\sigma}$ for some $N > 0$. Let τ'_i be the strategy where $i \in K \cup T$ begins by tossing N random coins, it then communicates the outcome of the coin tosses and its input to the adversary (using the communication scheme described in Section 6.1). Player i then plays τ_i using the outcome of the coin tosses whenever it needs to randomize. The scheduler σ'_E acts like σ_E except that it does not deliver any message that a player $j \in (K \cup T)$ sends with τ'_j that is not also sent with τ_j . (Note that for j to communicate its initial state and randomness to the scheduler does not actually require j to send messages to the scheduler. It just sends messages to itself, which do not have to be delivered.) By construction, we have that

$$\begin{aligned} & u'_i(\Gamma_d, (\vec{\sigma}_{-(K \cup T)}, \vec{\tau}_{(K \cup T)}), \sigma_E, \vec{x}_{(K \cup T)}) \\ = & u'_i(\Gamma_d, (\vec{\sigma}_{-(K \cup T)}, \vec{\tau}'_{(K \cup T)}), \sigma'_E, \vec{x}_{(K \cup T)}). \end{aligned}$$

We can view the adversary's strategy $(\vec{\tau}'_{(K \cup T)}, \sigma'_E)$ as a convex combination of (possibly infinitely many) deterministic strategies $(\vec{\tau}^*_{(K \cup T)}, \sigma_E^*)$. The construction of minimally-informative strategies guarantees that the number of honest players that terminate when running $(\Gamma_d, (\vec{\sigma}_{-(K \cup T)}, \vec{\tau}^*_{(K \cup T)}), \sigma_E^*, \vec{x}_{(K \cup T)})$ depends only on the scheduler and the inputs and randomization performed by players in $K \cup T$. Thus, given the input profile $\vec{x}_{(K \cup T)}$ of players in $K \cup T$, we can classify each of the deterministic strategies $(\vec{\tau}^*_{(K \cup T)}, \sigma_E^*)$ into the following three categories depending on how many honest players terminate (which, by our constraints on relaxed schedulers, must be the same in every history of $(\vec{\sigma}_{-(K \cup T)}, \vec{\tau}^*_{K \cup T}, \sigma_E^*)$):

- A. all honest players terminate;
- B. $n - 2t - 2k$ or more honest players do not terminate;
- C. at least one but fewer than $n - 2t - 2k$ honest players do not terminate.

Consider a scheduler σ''_e that acts just like σ'_E as long as the history is consistent with a history of $(\vec{\sigma}_{-(K \cup T)}, \vec{\tau}^*_{K \cup T}, \sigma'_E)$, until there comes a point when it is clear that some honest players will not terminate. If such a point comes, or if the history is inconsistent with a history of $(\vec{\sigma}_{-(K \cup T)}, \vec{\tau}^*_{K \cup T}, \sigma'_E)$, then σ''_e delivers all undelivered messages and from then on delivers all messages immediately. in more detail, σ''_e just like σ'_E until one of the following conditions holds:

- A player in $K \cup T$ does not communicate its initial state and the outcome of N coin tosses to the scheduler.

- The scheduler σ'_e can tell that what has happened thus far is inconsistent with the information sent by the players in $K \cup T$, assuming that they are using $\vec{\tau}'_{K \cup T}$ and the remaining players are using $\vec{\sigma}_{-(K \cup T)}$.
- It follows from the information sent by the players in $K \cup T$ that, with $(\sigma_{-(K \cup T)}, \tau_{K \cup T}, \sigma_E^*)$, some honest player will not terminate.
- All honest players terminate.

Note that, by construction, one of these one of these conditions must hold in every history. For if the mediator and honest players play a minimally-informative strategy, then the honest players do not randomize, and the mediator randomizes only with regard to the message it sends along with a STOP message. Moreover, an honest player terminates iff it gets a STOP message. Whether it gets one is determined by the scheduler's strategy, and the input of and randomization used by the players in $K \cup T$. If the players in $K \cup T$ use $\vec{\tau}_{K \cup T}$, then the scheduler can determine as soon as it has received the input and the coin tosses of the players in $K \cup T$ which honest players will terminate. Similarly, the scheduler can determine exactly when each honest player that terminates does so. In any case, once one of these conditions holds, the scheduler σ'_e delivers all of the messages not yet delivered, and from then on delivers messages immediately after they are sent. Thus, σ'_e is guaranteed to be standard.

Suppose that $(\vec{\tau}_{(K \cup T)}^*, \sigma_E^*)$ is a deterministic strategy in the support of $(\vec{\tau}'_{(K \cup T)}, \sigma'_E)$ that is in category A. Then, σ'_E and σ_e are indistinguishable when the players in $K \cup T$ play $\vec{\tau}_{(K \cup T)}^*$. Thus, if $\vec{\sigma} + \sigma_d$ is a minimally-informative ϵ -(k, t)-robust (resp., strong ϵ -(k, t)-robust) equilibrium, then there exists a value $\epsilon' < \epsilon$ such that

$$\begin{aligned}
& u'_i(\Gamma_d, (\vec{\sigma}_{-(K \cup T)}, \vec{\tau}_{(K \cup T)}^*), \sigma_E^*, \vec{x}_{(K \cup T)}) \\
&= u'_i(\Gamma_d, (\vec{\sigma}_{-(K \cup T)}, \vec{\tau}_{(K \cup T)}^*), \sigma'_E, \vec{x}_{(K \cup T)}) \\
&< u'_i(\Gamma_d, (\vec{\sigma}_{-T}, \vec{\tau}_T), \sigma_e, \vec{x}_T) + \epsilon' \quad [\text{by Proposition 6.9}]
\end{aligned}$$

for some (resp., for all) $i \in K$.

If $(\vec{\tau}_{(K \cup T)}^*, \sigma_E^*)$ is in category B, then again we have that if $\vec{\sigma} + \sigma_d$ is a minimally-informative ϵ -(k, t)-robust (resp., strong ϵ -(k, t)-robust) equilibrium, there exists a value $\epsilon' < \epsilon$ such that

$$\begin{aligned}
& u'_i(\Gamma_d, (\vec{\sigma}_{-(K \cup T)}, \vec{\tau}_{(K \cup T)}^*), \sigma_E^*, \vec{x}_{(K \cup T)}) \\
&< u'_i(\Gamma_d, \vec{\sigma}, \sigma'_e, \vec{x}_{(K \cup T)}) \quad [\text{by definition of } (2t + 2k)\text{-punishment strategy}] \\
&< u'_i(\Gamma_d, (\vec{\sigma}_{-T}, \vec{\tau}_T), \sigma_e, \vec{x}_T) + \epsilon' \quad [\text{by Proposition 6.8}]
\end{aligned}$$

for some (resp., for all) $i \in K$.

Since $(\vec{\sigma}_{-(K \cup T)}, \vec{\tau}_{(K \cup T)}, \sigma_E)$ ϵ -($K \cup T$)-($t + k + 1$)-coterminates, the probability that the strategy played by players in $K \cup T$ is in category C is at most ϵ . In this case, the payoff for each player is bounded by M . Using a compactness argument analogous to that of Proposition 6.8, there exists a value $\epsilon_0 < \epsilon$ such that $\epsilon' \leq \epsilon_0$ for all deterministic relaxed adversaries $(\vec{\tau}_{(K \cup T)}^*, \sigma_E^*)$ in the support of $(\vec{\tau}'_{(K \cup T)}, \sigma'_E)$ in categories A and B, and all inputs $\vec{x}_{(K \cup T)}$. So

$$\begin{aligned}
& u'_i(\Gamma_d, (\vec{\sigma}_{-(K \cup T)}, \vec{\tau}_{(K \cup T)}^*), \sigma_E^*, \vec{x}_{(K \cup T)}) \\
&< u'_i(\Gamma_d, (\vec{\sigma}_{-T}, \vec{\tau}_T), \sigma_e, \vec{x}_T) + \epsilon_0 + \epsilon M,
\end{aligned}$$

and therefore, if $\vec{\sigma} + \sigma_d$ is a minimally-informative ϵ -(k, t)-robust (resp., strong ϵ -(k, t)-robust) equilibrium, then

$$\begin{aligned} & u'_i(\Gamma_d, (\vec{\sigma}_{-(K \cup T)}, \vec{\tau}_{(K \cup T)}), \sigma_E, \vec{x}_{(K \cup T)}) \\ & < u'_i(\Gamma_d, (\vec{\sigma}_{-T}, \vec{\tau}_T), \sigma_e, \vec{x}_T) + \epsilon_0 + \epsilon M \end{aligned}$$

for all inputs $\vec{x}_{(K \cup T)}$, all standard schedulers σ_e , all strategies $\vec{\tau}_{(K \cup T)}$, and for some (resp., for all) $i \notin T$.

It remains to prove the result in the case that A is not N -bounded with respect to $\vec{\sigma}$ for some N . Fix $\epsilon' > 0$ and let A'' be an N -bounded ϵ' -approximation of A with respect to $\vec{\sigma}$ given by Lemma 6.14. Then, by the previous argument

$$\begin{aligned} & u'_i(\Gamma_d, (\vec{\sigma}_{-(K \cup T)}, \vec{\tau}_{(K \cup T)}''), \sigma_E'', \vec{x}_{(K \cup T)}) \\ & < u'_i(\Gamma_d, (\vec{\sigma}_{-T}, \vec{\tau}_T), \sigma_e, \vec{x}_T) + \epsilon_0 + \epsilon M, \end{aligned}$$

and by Lemma 6.14 it follows that

$$\begin{aligned} & u'_i(\Gamma_d, (\vec{\sigma}_{-(K \cup T)}, \vec{\tau}_{(K \cup T)}), \sigma_E, \vec{x}_{(K \cup T)}) \\ & < u'_i(\Gamma_d, (\vec{\sigma}_{-T}, \vec{\tau}_T), \sigma_e, \vec{x}_T) + \epsilon_0 + \epsilon M + \epsilon' M \end{aligned}$$

for all input profiles $\vec{x}_{(K \cup T)}$, standard schedulers σ_e , and strategy profiles $\vec{\tau}_{(K \cup T)}$, and some (resp., all) $i \notin T$. Since this inequality holds for all $\epsilon' > 0$, the result follows. $\square$

An analogous argument can be used if we have an (k, t) -robust equilibrium (not just an ϵ -(k, t)-robust equilibrium):

Proposition 6.15. *If $\vec{\sigma} + \sigma_d$ is a minimally-informative (k, t) -robust equilibrium in a mediator game $\Gamma_d(u'_i)$ for which a $(k + t)$ -punishment strategy exists, σ_E is a relaxed scheduler, T and K are disjoint sets of players with $|T| \leq t$ and $1 \leq |K| \leq k$, and $\vec{\tau}_{(K \cup T)}$ is a strategy profile for the players in $K \cup T$ such that $(\vec{\sigma}_{-(K \cup T)}, \vec{\tau}_{(K \cup T)}, \sigma_E)$ $(K \cup T)$ -coterminates, then there exists a (standard) scheduler σ_e such that for all input profiles $\vec{x}$ and all $i \notin T$,*

$$\begin{aligned} & u'_i(\Gamma_d, (\vec{\sigma}_{-(K \cup T)}, \vec{\tau}_{(K \cup T)}), \sigma_d, \sigma_E, \vec{x}_{(K \cup T)}) \\ & \leq u'_i(\Gamma_d, (\vec{\sigma}_{-T}, \vec{\tau}_T), \sigma_d, \sigma_e, \vec{x}_T). \end{aligned}$$

Returning to the proof of Theorem 4.4, we can now prove (strong) (k, t) -robustness. Let $\Gamma_d(\vec{u}')$ be a utility variant of Γ_d such that $\vec{\sigma} + \vec{\sigma}_d$ is a (k, t) -robust equilibrium of $\Gamma_d(\vec{u}')$, let σ_e be a scheduler in $\Gamma_{CT}(\vec{u})$, and let K and T be disjoint subsets of players with $1 \leq |K| \leq k$ and $|T| \leq t$, respectively, such that $3k + 4t < n$. Let $\vec{\tau}_K$ and $\vec{\tau}_T$ be strategy profiles for players in K and T , respectively. By Theorem 5.3 and Proposition 6.1, there exist a function H_{σ_e} and a relaxed scheduler σ_E in the mediator game such that

$$\begin{aligned} & u'_i(\Gamma_{CT}(\vec{u}'), (\vec{\sigma}_{CT})_{-(K \cup T)}, \vec{\tau}_K, \vec{\tau}_T), \sigma_e, \vec{x}_{(K \cup T)}) \\ & = u'_i(\Gamma_d(\vec{u}'), (\vec{\sigma}_{-(K \cup T)}, H_{\sigma_e}(\vec{\tau}_K), H_{\sigma_e}(\vec{\tau}_T)), \sigma_d, \sigma_E, \vec{x}_{(K \cup T)}) \end{aligned}$$

for all $i \notin T$ and all input profiles $\vec{x}$. We can assume without loss of generality that $(\vec{\sigma}, \sigma_d)$ is minimally informative. Thus, by Proposition 6.15, there exists a standard scheduler σ'_e such that

$$\begin{aligned} & u'_i(\Gamma_d(\vec{u}'), (\vec{\sigma}_{-(K \cup T)}, H_{\sigma_e}(\vec{\tau}_T), H_{\sigma_e}(\vec{\tau}_K)), \sigma_d, \sigma_E, \vec{x}_{(K \cup T)}) \\ & \leq u'_i(\Gamma_d(\vec{u}'), (\vec{\sigma}_{-T}, H_{\sigma_e}(\vec{\tau}_T)), \sigma_d, \sigma'_e, \vec{x}_{(K \cup T)}). \end{aligned}$$

Finally, by Theorem 5.3, if $\vec{\sigma}$ is (k, t) -robust (resp., strongly (k, t) -robust), then there exists a standard scheduler σ_e'' such that

$$\begin{aligned} & u'_i(\Gamma_{CT}(\vec{u}'), ((\vec{\sigma}_{CT})_{-T}, \vec{\tau}_T), \sigma_e, \vec{x}_T) \\ &= u'_i(\Gamma_d(\vec{u}'), (\vec{\sigma}_{-T}, H_{\sigma_e}(\vec{\tau}_T)), \sigma_d, \sigma_e'', \vec{x}_T) \quad [\text{by Theorem 5.3}] \\ &= u'_i(\Gamma_d(\vec{u}'), (\vec{\sigma}_{-T}, H_{\sigma_e}(\vec{\tau}_T)), \sigma_d, \sigma_e', \vec{x}_T) \quad [\text{by Corollary 6.5}] \end{aligned}$$

for some $i \in K$ (resp., for all $i \in K$). Therefore,

$$\begin{aligned} & u'_i(\Gamma_{CT}(\vec{u}'), ((\vec{\sigma}_{CT})_{-(K \cup T)}, \vec{\tau}_K, \vec{\tau}_T), \sigma_e, \vec{x}_{(K \cup T)}) \\ &\leq u'_i(\Gamma_{CT}(\vec{u}'), ((\vec{\sigma}_{CT})_{-T}, \vec{\tau}_T), \sigma_e, \vec{x}_T), \end{aligned}$$

as desired. $\square$

We remark that, with a little more effort, we can show that the minimally-informative strategy profile $f(\vec{\sigma} + \sigma_d)$ that implements $\vec{\sigma} + \sigma_d$ is actually a $(t + k)$ -bisimulation of $\vec{\sigma} + \sigma_d$. Moreover, the strategy profile that implements $f(\vec{\sigma} + \sigma_d)$ in the cheap-talk game preserves all the properties of the cheap-talk equilibrium in Theorem 4.4. Thus, under the conditions of Theorem 4.4, we can get a strategy profile in the cheap-talk game that $(t + k, t)$ -bisimulates a strategy profile in the mediator game.

6.6 Proof of Theorem 4.5

To prove Theorem 4.5, we use an analogous strategy to that used to prove Theorem 4.4, using Theorem 5.4 instead of Theorem 5.3. The same argument as that used in the proof Theorem 4.2 shows that for all $\epsilon' \in (0, 1]$ there exists a protocol $\vec{\sigma}_{CT}$ that ϵ' -implements $\vec{\sigma} + \sigma_d$ and that $\vec{\sigma}_{CT}$ is (ϵ, t) -immune.

To prove (strong) ϵ -(k, t)-robustness, fix an adversary $A = (\vec{\tau}_K, \vec{\tau}_T, \sigma_e)$ for subsets K, T such that $1 \leq |K| \leq k$, $|T| \leq t$ and $K \cap T = \emptyset$. By Theorem 5.4 and Proposition 6.2, there exists a function H_{σ_e} from strategies to strategies and a relaxed scheduler σ_E such that, for all input profiles $\vec{x}$,

$$\begin{aligned} & u_i(\Gamma_{CT}, ((\vec{\sigma}_{CT})_{-(K \cup T)}, \vec{\tau}_K, \vec{\tau}_T), \sigma_e, \vec{x}_{(K \cup T)}) \\ &< u_i(\Gamma_d, (\vec{\sigma}_{-(K \cup T)}, H_{\sigma_e}(\vec{\tau}_K), H_{\sigma_e}(\vec{\tau}_T)), \sigma_d, \sigma_E, \vec{x}_{(K \cup T)}) + \epsilon' M \end{aligned}$$

for all $i \in K$.

By Theorem 5.4, there exists a standard scheduler σ_e' such that

$$\begin{aligned} & u_i(\Gamma_{CT}, ((\vec{\sigma}_{CT})_{-T}, \vec{\tau}_T), \sigma_e, \vec{x}_T) \\ &> u_i(\Gamma_d, (\vec{\sigma}_{-T}, H_{\sigma_e}(\vec{\tau}_T)), \sigma_d, \sigma_e', \vec{x}_T) - \epsilon' M. \end{aligned}$$

Thus, there exists some $\epsilon_0 < \epsilon$ such that if $\vec{\sigma}_d$ is (k, t) -robust (resp., strongly (k, t) -robust), then

$$\begin{aligned} & u_i(\Gamma_{CT}, ((\vec{\sigma}_{CT})_{-(K \cup T)}, \vec{\tau}_K, \vec{\tau}_T), \sigma_e, \vec{x}_{(K \cup T)}) \\ &< u_i(\Gamma_d, (\vec{\sigma}_{-(K \cup T)}, H_{\sigma_e}(\vec{\tau}_K), H_{\sigma_e}(\vec{\tau}_T)), \sigma_d, \sigma_E, \vec{x}_{(K \cup T)}) + \epsilon' M \\ &< u_i(\Gamma_d, (\vec{\sigma}_{-T}, H_{\sigma_e}(\vec{\tau}_T)), \sigma_d, \sigma_e', \vec{x}_T) + \epsilon_0 + 2\epsilon' M \quad [\text{by Proposition 6.12}] \\ &< u_i(\Gamma_{CT}, ((\vec{\sigma}_{CT})_{-T}, \vec{\tau}_T), \sigma_e, \vec{x}_T) + \epsilon_0 + 3\epsilon' M \end{aligned}$$

for some $i \in K$ (resp., for all $i \in K$). Therefore, taking $\epsilon' = (\epsilon - \epsilon_0)/3M$, we have that $\vec{\sigma}_{CT}$ is a ϵ -(k, t)-robust equilibrium (resp., strongly (k, t) -robust equilibrium) for Γ_{CT} .

Again, as was the case for Theorem 4.4, with a little more effort we can show that under the conditions of Theorem 4.5, we can get a strategy profile in the cheap-talk game that ϵ -($t + k, t$)-bisimulates a strategy profile in the mediator game.

7 Conclusion

We have extended the results of ADGH on implementing mediators to the asynchronous setting. This setting raises a number of new subtleties, particularly regarding how to define utilities in a game where players do not terminate. Since many real-world applications are asynchronous, and thinking in terms of mediators in this setting provides much simpler approach to designing efficient mechanisms, having a “compiler” that can translate solutions with a mediator to one without a mediator can be quite useful in principle.

There are still a number of questions that remain open. The most obvious ones involve lower bounds. Lower bounds that match the upper bounds of ADGH in the synchronous setting were proved by Abraham, Dolev, and Halpern [2008]. Can we provide analogous lower bounds here? In addition, we have considered only non-cryptographic setting; ADGH also provided bounds for the setting where players were polynomially-bounded and could use cryptographic tools. To what extent do things change in this setting in the asynchronous case?

References

- I. Abraham, D. Dolev, I. Geffner, and J. Y. Halpern. 2019. Implementing mediators with asynchronous cheap talk. (2019). Available at <http://arxiv.org/abs/1806.01214>.
- I. Abraham, D. Dolev, R. Gonen, and J. Y. Halpern. 2006. Distributed computing meets game theory: robust mechanisms for rational secret sharing and multiparty computation. In *Proc. 25th ACM Symposium on Principles of Distributed Computing*. 53–62.
- I. Abraham, D. Dolev, and J. Y. Halpern. 2008. Lower bounds on implementing robust and resilient mediators. In *Fifth Theory of Cryptography Conference*. 302–319.
- E. Adar and B. Huberman. 2000. Free riding on Gnutella. *First Monday* 5, 10 (2000).
- R. J. Aumann and S. Hart. 2003. Long cheap talk. *Econometrica* 71, 6 (2003), 1619–1660.
- M. Ben-Or, R. Canetti, and O. Goldreich. 1993. Asynchronous secure computation. In *STOC '93: Proceedings of the 25 Annual ACM Symposium on Theory of Computing*. 52–61.
- M. Ben-Or, S. Goldwasser, and A. Wigderson. 1988. Completeness theorems for non-cryptographic fault-tolerant distributed computation. In *Proc. 20th ACM Symp. Theory of Computing*. 1–10.
- M. Ben-Or, B. Kelmer, and T. Rabin. 1994. Asynchronous secure computations with optimal resilience (extended abstract). In *Proc. 13th ACM Symp. Principles of Distributed Computing*. 183–192.
- E. Ben-Porath. 2003. Cheap talk in games with incomplete information. *Journal of Economic Theory* 108, 1 (2003), 45–71.

- S. Even, O. Goldreich, and A. Lempel. 1985. A randomized protocol for signing contracts. *Commun. ACM* 28, 6 (1985), 637–647. <https://doi.org/10.1145/3812.3818>
- I. Geffner and J. Y. Halpern. 2018. Stronger security guarantees for multiparty computation. (2018). Available at <http://arxiv.org>.
- O. Goldreich, S. Micali, and A. Wigderson. 1987. How to play any mental game. In *Proc. 19th ACM Symp. Theory of Computing*. 218–229.
- J. Y. Halpern and M. R. Tuttle. 1993. Knowledge, probability, and adversaries. 40, 4 (1993), 917–962.
- A. Shamir, R. L. Rivest, and L. Adelman. 1981. Mental poker. In *The Mathematical Gardner*, D. A. Klarner (Ed.). Prindle, Weber, and Schmidt, Boston, MA, 37–43.
- A. Yao. 1982. Protocols for secure computation (extended abstract). In *Proc. 23rd IEEE Symp. Foundations of Computer Science*. 160–164.