



Munich Personal RePEc Archive

Communication, correlation and cheap-talk in games with public information

Heller, Yuval and Solan, Eilon and Tomala, Tristan

Tel-Aviv University

1 June 2010

Online at <https://mpra.ub.uni-muenchen.de/25895/>
MPRA Paper No. 25895, posted 18 Oct 2010 15:01 UTC

Communication, correlation and cheap-talk in games with public information

Yuval Heller

The School of Mathematical Sciences, Tel Aviv University Tel Aviv, 69978 Israel.

Eilon Solan

The School of Mathematical Sciences, Tel Aviv University Tel Aviv, 69978 Israel.

Tristan Tomala

*HEC Paris, Economics and Decision Sciences Department, 78351 Jouy en Josas,
France.*

Abstract

This paper studies extensive form games with public information where all players have the same information at each point in time. We prove that when there are at least three players, all communication equilibrium payoffs can be obtained by unmediated cheap-talk procedures. The result encompasses repeated games and stochastic games.

1 Introduction

Game theory models rational agents as selfish players who take actions independently of each other. In reality, players' decisions often depend on correlated external events (sunspots) and players may exchange messages before taking decisions. The correlation of actions was formalized in the seminal work of Aumann (1974) who showed that correlated actions may achieve (Pareto-)better outcomes. Aumann's correlated equilibrium requires a centralized

Email addresses: helleryu@post.tau.ac.il (Yuval Heller), eilons@post.tau.ac.il (Eilon Solan), tomala@hec.fr (Tristan Tomala).

and trustworthy mediator, whose existence is generally a demanding assumption. An alternative model that allows players to correlate their actions involves cheap-talk, where players communicate directly with each other in a non-binding and costless way (see, e.g., Crawford and Sobel, 1982; Farrell and Rabin, 1996). Many papers study the implementation of correlated equilibria by such decentralized communication (Abraham et al, 2006, 2008; Barany, 1992; Ben-Porath, 1998, 2003; Gerardi, 2004).

Most of the literature on cheap-talk concerns static games, with or without complete information. The very nature of sequential games allows for various types of correlation mechanisms (Forges, 1986; Myerson, 1986): the mediator can send messages before the beginning of the game, send additional messages during the play and receive messages from the players. A mediator who only sends pre-play messages gives rise to a *normal form* correlated equilibrium. If the mediator sends further messages at each stage, it gives rise to an *extensive form* correlated equilibrium. When the mediator exchanges messages with the players all through the game, the corresponding equilibrium concept is termed *communication* equilibrium (see Forges, 1986, for this classification). This latter equilibrium concept encompasses all others. A natural question is then whether in sequential games, any communication equilibrium can be implemented by using cheap-talk, without the help of a mediator. Another issue is the implementation of communication by pre-play procedures (mediated or not). Indeed, in some applied settings, players may be able to communicate only before the start of the game; in others, communication during the game may be costly and insecure. For example, in the midst of a military action, communication among units of the same army may be insecure or even impossible. On the stock market, traders receiving an important piece of news need to act quickly, and every minute devoted to communication may have dramatic effects on performance (see Heller, 2010b). The aim of the present paper is to study the implementation of communication equilibria by cheap-talk and/or pre-play procedures.

We consider extensive form games with public information (Dubey and Kaneko, 1984; Osborne and Rubinstein, 1994), where at each point of time, all players have the same information about the past history of the game. The length of the game is possibly infinite. These games include repeated games and, more generally, stochastic games with perfect monitoring of actions, where the players have symmetric information on the state variable.

Our first result (Theorem 9) shows that any communication equilibrium can be implemented by a pre-play correlation device (a mediator sends messages to the players before the game starts) complemented by a simple cheap-talk mechanism, where every player sends a single public message before each stage.

Our second result (Theorem 10) shows that when there are at least three players, one can

replace the mediator by a cheap-talk phase that takes place before the game starts. As a consequence, communication equilibria are implementable by cheap-talk procedures. The cheap-talk mechanisms we use have two alternative forms. In the first form, the players perform a long cheap-talk phase before the game starts, thereby exchanging many private messages. During the play, short cheap-talk phases are performed whereby each player sends a single public message. In the second form, the players perform cheap-talk phases before the game starts and at every stage along the play. The length of each cheap-talk phase is random, but the expected number of messages sent at each phase is finite.

We now discuss the main ingredients of the proofs of these results. To prove Theorem 9, we first strengthen the result of Solan (2001) and show that in games with public information, communication ε -equilibria are equivalent to extensive form correlated ε -equilibria. I.e., it is not essential to assume that the mediator receives messages from the players or observes the actual history of the game. Thanks to the revelation principle (see Forges (1986) and Myerson (1986)), any communication equilibrium can be implemented by a device which observes the history and sends recommendations that are obediently followed by the players. If the mediator does not observe the history nor receives messages, it is enough to let it send lists of history-dependent recommendations, and to let players coordinate on the messages relevant to the actual history. Second, we let the mediator act only at the pre-play stage. We use authentication schemes à la Rabin and Ben-Or (1989) to let the device send to each player encrypted recommended actions for the whole game. The encoding keys are told to another player. At each stage of the game, players simultaneously broadcast the encoding keys. The authentication properties of the schemes of Rabin and Ben-Or enable all players to know whether a broadcasted key is genuine or not.

To prove Theorem 10, we rely on the secure multiparty computation protocols of Rabin and Ben-Or (1989), and of Ben-Or et al. (1988). These cheap-talk protocols allow players to jointly compute *outputs* which are polynomial functions of the profile of private inputs of players. The computation is secure in that player i learns his own output without getting any information on the inputs and outputs of the other players. These protocols have been used for cheap-talk implementation of correlated equilibria in one-stage games in Abraham et al. (2006, 2008) and Heller (2010a). The novelty of the present paper is the adaptation of these protocols for the implementation of communication equilibria in multi-stage games.

Stochastic games are a special kind of games with public information, where the players perfectly observe the state variable and the action profile. Vieille (2000a, 2000b) proved that any two-player undiscounted stochastic game (with a finite number of states) admits an equilibrium payoff (without any communication). Whether this holds true for stochastic

games with more than two players is an open problem. Solan and Vieille (2002) proved that any undiscounted n -player stochastic game (with a finite number of states) admits an extensive-form correlated equilibrium. Our results yield the following corollary: Every undiscounted n -player stochastic game (with a finite number of states) admits a cheap-talk equilibrium payoff, i.e., a communication equilibrium payoff that involves only cheap-talk, with one of the two cheap-talk mechanisms described above.

The paper is organized as follows. The model and the results are described in Section 2. The proof of the first main result is given in Section 3, and the proof of the second main result is given in Section 4. We conclude in Section 5.

2 Model

2.1 Games with Public information

We study a class of extensive form games, henceforth called games with public information, where there is a timing structure, and at each point of time, all players have the same information about the past history of the game. These are multi-stage games, where at each stage, the moves of each player and of chance are publicly disclosed.¹ The game played at each stage can be history dependent. This class of games has been described in the literature as extensive games with perfect information and simultaneous moves (see Osborne and Rubinstein, 1994, page 102, based on Dubey and Kaneko, 1984), or as multi-stage games (see Forges, 1986). Let us define such games formally, following Osborne and Rubinstein.

A *game with public information* is a tuple $G = \langle I, H, P, A, f, (u^i) \rangle$ where:

- I is a finite set of players.
 - H is a set of sequences, finite or infinite, called histories. A history is denoted $h = (a_k)_{k=1, \dots, K}$ where $K \in \mathbb{N} \cup \{+\infty\}$ is the length of h . The following three properties are assumed:
 - ◊ The empty sequence \emptyset is a member of H .
 - ◊ A prefix of a history is a history: If $(a_k)_{k=1}^K \in H$ and $L < K$ then $(a_k)_{k=1}^L \in H$.
 - ◊ If all prefixes of an infinite sequence $(a_k)_{k=1}^\infty$ are histories, then so is the infinite sequence. That is, if $(a_k)_{k=1}^L \in H$ for every positive integer L then $(a_k)_{k=1}^\infty \in H$.
- A history $(a_k)_{k=1}^K \in H$ is *terminal* if it is infinite, or if there is no a_{K+1} such that $(a_k)_{k=1}^{K+1} \in H$. The set of terminal histories is denoted Z .

¹ All of our results hold if players have symmetric partial information about chance moves.

- P is a mapping that assigns to each nonterminal history h the set of players $P(h) \subseteq I$ that have to take an action after history h . If $P(h) = \emptyset$ then there is a chance move after history h .
- A is a mapping that assigns to every nonterminal history h such that $P(h) \neq \emptyset$, and to every player $i \in P(h)$, a finite set $A^i(h)$ of actions available to player i after that history. Let $A(h)$ be the set of available action-profiles at h : $A(h) = \times_{i \in P(h)} A^i(h)$. If $P(h) = \emptyset$ for some nonterminal history h , then $A(h)$ is the finite set of chance moves at the history h .

The set of histories H , and the function A satisfy the following property. For every nonterminal history h : $a \in A(h) \Leftrightarrow (h, a) \in H$. That is, a history $h = (a_k)_{k=1, \dots, K}$ is a sequence of action profiles where the components of a_k are the actions taken by players $i \in P\left((a_l)_{l=1}^{k-1}\right)$ or by chance (if $P\left((a_l)_{l=1}^{k-1}\right) = \emptyset$).

- f is a mapping that assigns to every nonterminal history h such that $P(h) = \emptyset$, a probability distribution $f(\cdot | h)$ over chance moves $A(h)$. That is, when chance has to move after a nonterminal history h , an action $a \in A(h)$ is chosen according to the probability distribution $f(\cdot | h)$.
- For each player $i \in I$, $u^i : Z \rightarrow [0, 1]$ is the payoff function of player i defined over terminal histories. This function is assumed to be measurable with respect to the product σ -algebra on H ; the σ -algebra over each finite set $A(h)$ is the discrete σ -algebra.

The game unfolds as follows. The empty history is the starting point of the game. Players in $P(\emptyset)$ choose actions simultaneously (if $P(\emptyset) = \emptyset$, then chance chooses an action according to the distribution $f(\cdot | \emptyset)$). Given the chosen action profile a , players in $P(a)$ choose actions at the next stage and so on until a terminal history z is reached (recall that histories can be infinite and that an infinite history is terminal). Each player $i \in I$ receives the payoff $u_i(z)$.

Games with public information encompass extensive form games without information sets, repeated games with perfect monitoring of actions where all players move at each stage, and more generally, stochastic games with perfect monitoring of state and actions, where the current game depends on a parameter that evolves according to the moves of the players and of chance. In fact, any game with public information can be represented as a stochastic game, where H is the state space and the law of motion is the one described above by the data of P and f .

2.2 Communication and correlated equilibria

Since the seminal work of Aumann (1974) on correlated equilibria, various solution concepts extending Nash equilibria have been proposed to account for possibilities of costless communication between the players. We present now the main solution concepts, following Forges (1986) and Myerson (1986).

A *communication device* is an agent that exchanges messages with the players between game stages. This models a trustworthy mediator, which helps the players communicate and correlate their actions. It specifies spaces of messages that the device sends to the players, spaces of messages that the device receives from the players, and the rule according to which the device sends messages.

Formally, let G be a game with public information. A *communication device* is a tuple $D = \left((S^i(h))_{i \in I, h \in H \setminus Z}, (R^i(h))_{i \in I, h \in H \setminus Z}, \mu \right)$ where:

- For each player i , $S^i(h)$ is a measurable set of signals that the device can send to player i after history h , and $R^i(h)$ is a measurable set of messages that the device can receive from player i after history h .

An *extended history* is a triple (h, s, r) where $h = (a_k)_{k=1}^K$ is a nonterminal history of the game, $s = (s_k)_{k=1}^K$, and $r = (r_k)_{k=1}^K$ are feasible histories of messages, i.e., for each $n < K$, $s_{n+1} \in S(h_n) := \times_{i \in I} S^i(h_n)$, and $r_{n+1} \in R(h_n) := \times_{i \in I} R^i(h_n)$, with $h_n := (a_k)_{k \leq n}$.

- μ is a transition probability that maps extended histories to probability distributions over signals sent to the players: $\mu(\cdot | h, s, r) \in \Delta(S(h))$ is a probability distribution over $S(h)$.

Given a communication device D , *the game extended by D* , noted $G(D)$, unfolds as follows. After each extended history $(h, s, r) = \left((a_k)_{k=1}^K, (s_k)_{k=1}^K, (r_k)_{k=1}^K \right)$:

- (1) The device chooses a profile of signals $s_{K+1} = (s_{K+1}^i) \in S(h)$ according to $\mu(h, s, r)$. Each player i is privately informed of s_{K+1}^i .
- (2) Each player $i \in P(h)$ chooses an action a_{K+1}^i in $A^i(h)$ (if $P(h) = \emptyset$ then chance's move $a_{K+1} \in A(h)$ is randomly chosen according to $f(\cdot | h)$). The selected action profile (or chance's move) a_{K+1} is publicly announced.
- (3) Each player $i \in I$ sends a private message $r_{K+1}^i \in R^i(h)$ to the device.

Remark 1 The definition of a communication device adopted here is called a *general communication device* in Solan (2001), since in the original definition of Forges (1986), the device does not observe the history of the game. However, in games with public information, this is the same concept. Indeed, at each stage K , the device may send to each player a vector

of messages, one for each possible history of length K . The recommendations for a given history depend on past recommendations given along this history. In other words, the device simulates in parallel all possible executions of the game and proceeds in each instance as if it were the actual game. Since histories are common knowledge, all players know which message to take note of, and the messages associated to unrealized histories are irrelevant to them.

Throughout the paper, we assume that players have perfect recall and use behavior strategies. A behavior strategy of player i in $G(D)$ is a function $y^i = (x^i, m^i)$ mapping the extended histories of observations of player i to probability distributions over actions or messages chosen by player i . That is, let $(h, s, r) = ((a_k)_{k=1}^K, (s_k)_{k=1}^K, (r_k)_{k=1}^K)$ be an extended history. At stage K , player i has observed $(h, (s_k^i)_{k=1}^K, (r_k^i)_{k=1}^K)$ and receives the new signal s_{K+1}^i . Then $x^i(h, (s_k^i)_{k=1}^{K+1}, (r_k^i)_{k=1}^K)$ is the probability distribution over $A^i(h)$ used by player i for choosing his new action (whenever $i \in P(h)$). After actions have been chosen, player i has observed $(h, a_{K+1}, (s_k^i)_{k=1}^{K+1}, (r_k^i)_{k=1}^K)$ and chooses a new message r_{k+1}^i according to a distribution over $R^i(h)$ denoted $m^i(h, a_{K+1}, (s_k^i)_{k=1}^{K+1}, (r_k^i)_{k=1}^K)$.

We denote by $\gamma_D^i(y) = E_y(u^i(z))$ the expected payoff of player i with respect to the probability distribution induced by the correlation device D and the strategy profile y over terminal histories. For $\varepsilon \geq 0$, a strategy profile y is an ε -Nash equilibrium of the extended game $G(D)$ if for every player $i \in I$ and every strategy \hat{y}^i of player i : $\gamma_D^i(y) \geq \gamma_D^i(y^{-i}, \hat{y}^i) - \varepsilon$, where $-i$ denotes $I \setminus \{i\}$ and $y^{-i} = (y^j)_{j \neq i}$.

Definition 2 Let G be a game with public information and $\varepsilon \geq 0$. A *communication ε -equilibrium* of G is a communication device D and an ε -Nash equilibrium of $G(D)$. A payoff vector $g \in R^I$ induced by a communication ε -equilibrium is a *communication ε -equilibrium payoff*. A payoff vector $g \in R^I$ is a *communication equilibrium payoff* if it is the limit of communication ε -equilibrium payoffs as $\varepsilon > 0$ goes to 0.

Remark 3 A communication 0-equilibrium payoff is a communication equilibrium payoff. The converse need not be true. It is possible that communication 0-equilibria do not exist whereas communication equilibrium payoffs do (see, e.g., the ‘‘Big Match’’ in Blackwell and Ferguson, 1968).

Special classes of communication devices are the following:

- A communication device D is *canonical* if the mediator does not receive inputs from the players ($R^i(h)$ is a singleton for all i, h), and the signal it sends to each player i is a recommended action that player i should play at the next stage: $S^i(h) = A^i(h)$ if $i \in P(h)$, a singleton otherwise.
- A communication device is *autonomous* if the mediator does not receive inputs from the

players ($R^i(h)$ is a singleton) and does not observe the history of the game ($S^i(h)$ and $\mu(\cdot|h, r, s)$ depend on h only through its length K).

- A communication device is a *pre-play correlation device* if it only sends messages before the beginning of the game, that is $S^i(h)$ is a singleton unless $h = \emptyset$.

When the communication device D is canonical, one strategy that is available to each player i is the *obedient* strategy $\hat{\sigma}^i$ that follows the device's recommendation. For $\varepsilon \geq 0$, $(D, (\hat{\sigma}^i)_{i \in I})$ is a *canonical communication ε -equilibrium* if D is a canonical communication device and the obedient strategy profile $\hat{\sigma} = (\hat{\sigma}^i)_{i \in I}$ is an ε -equilibrium of $G(D)$. For $\varepsilon \geq 0$, an *extensive form correlated ε -equilibrium* (*correlated ε -equilibrium*) of G is a communication ε -equilibrium induced by an autonomous (pre-play correlation) device.

Remark 4 A revelation principle applies to communication equilibria (see Forges, 1986; Myerson, 1986). That is, any communication ε -equilibrium is equivalent to a *canonical* communication ε -equilibrium where the device recommends actions to the players, at equilibrium each player actually plays the recommended action, and then players faithfully report their incremental information to the device. Here, the reports of the players consist in announcing the newly played action profile, which is superfluous since the device observes the history. It is thus without loss of generality to assume that the players do not send messages.

This discussion leads to the following proposition.

Proposition 5 *Let G be a game with public information. For every $\varepsilon \geq 0$, every communication ε -equilibrium is equivalent to: (1) a canonical communication ε -equilibrium, and (2) an extensive form correlated ε -equilibrium.*

A similar result is proved in Solan (2001), who shows that for games with public information and general action spaces, communication and extensive form correlated equilibrium payoffs coincide. Proposition 5 is slightly stronger: every communication ε -equilibrium can be exactly replicated by an autonomous (or canonical) device. The first part of the proposition directly follows from the revelation principle. The proof of the second part of the proposition is a building block of the proof of Theorem 9, and is given in Section 3 for the sake of completeness.

Proposition 5 is specific to games with public information. For instance, in repeated games with imperfect private monitoring of actions, communication and extensive form correlated equilibria are not equivalent (see Renault and Tomala, 2004). It is known, however, that pre-play correlated equilibria are not equivalent to extensive form correlated equilibria, even in games with public information, see Forges, 1986.

2.3 Cheap-Talk

Cheap-talk is a particular form of communication where players can freely and costlessly exchange messages without any mediation.² In our cheap-talk model, we assume that each player is able to send a private message to any other player (and no other player can intercept this message), and that each player is able to broadcast public announcements.³ In addition we assume that the identity of the sender of each message is certifiable.

Definition 6 • A cheap-talk phase specifies a finite message space M containing a null message \diamond , and consists of (possibly infinitely many) rounds of communication. In each round n , each player i can send simultaneous private and public messages (i.e. send a private message to each player j and/or broadcast a message).

- A cheap-talk extension G^* of a game with perfect information G is a game in extensive form where, after each non-terminal history h , a cheap-talk phase is played with a history dependent message space $M(h)$.
- A cheap-talk ε -equilibrium payoff of G is an ε -equilibrium payoff of a cheap-talk extension G^* of G . A cheap-talk equilibrium payoff is the limit of ε -equilibrium payoffs as $\varepsilon > 0$ goes to 0.

Cheap-talk extensions are particular kinds of communication devices, and consequently a cheap-talk ε -equilibrium is a communication ε -equilibrium. Though a cheap-talk phase can have infinitely many rounds, in most of our constructions, the number of communication rounds is either finite, or has finite expectation. Denote by $s_{k,n}^{i,j}$ (resp. $s_{k,n}^{i,I}$) the private (resp. public) message that player i sends to player j (resp. broadcasts) at the n -th round of the k -th cheap-talk phase. An i -information set in a cheap-talk extension after the N -th round of the K -th cheap-talk phase is (h, s^i, N) where $h = (a_k)_{k=1}^K$ is a history of the game, $s^i = (s_k^i)_{k=1}^{K+1}$ is the history of messages that player i sent or received in past cheap-talk phases $(1, \dots, K)$, and in the current phase $(K+1)$ until round N (N may be finite or infinite). That is, for each $k < K + 1$, $s_k^i = \left(s_{k,n}^{i,j}, s_{k,n}^{j,i}, s_{k,n}^{j,I} \right)_{j \in I, n \geq 1}$ is the sequence of messages player i sent or received in the k 'th cheap-talk phase, and $s_{K+1}^i = \left(s_{K+1,n}^{i,j}, s_{K+1,n}^{j,i}, s_{K+1,n}^{j,I} \right)_{j \in I, 1 \leq n \leq N}$ is the sequence of messages that player i sent or received in the first N rounds of the $K + 1$ 'th cheap-talk phase. A behavior strategy of player i in G^* is denoted by $y^i = \left(x^i, (m_N^i)_{N \geq 1} \right)$ and

² See Farrel and Rabin (1996) for a nontechnical introduction to some of the main issues of cheap-talk.

³ When there are four or more players, the constructions may be adapted to use only 2-player private channels. Cryptographic assumptions (players are computationally restricted and “one-way” functions exist) are needed to adapt the constructions to use only public announcements, see Abraham et al., 2006, 2008; and Urbano and Vila, 2002).

maps i -information sets to distributions over actions and messages. For an i -information set (h, s^i, ∞) (∞ indicates the end of a cheap-talk phase), we denote $x^i(h, s^i, \infty)$ the probability distribution of the next action chosen of player i . For each i -information set (h, s^i, N) where N is finite, we denote $m_n^i(h, s^i, N)$ the distribution of messages sent by player i (a private message for each player j and a public message) at the $(N + 1)$ -th round of the cheap-talk phase that follows history h .

An extended history in a cheap-talk extension after the N -th round of the K -th cheap-talk phase is a profile of i -information sets for each player: $(h, s, N) = (h, s^i, N)_{i \in I}$. Let y be a strategy profile in G^* and $(h, s) := (h, s, 0)$ an extended history at the beginning of the cheap-talk phase that follows history h . The *length of the cheap-talk phase* that follows (h, s) is a random variable denoted $l_y(h, s)$. That is, $l_y(h, s)$ is the minimal n_0 , such that for each round $n \geq n_0$, all messages that are sent by the players are equal to \diamond . If there exists no such n_0 , $l_y(h, s) = \infty$.

Definition 7 A strategy profile y in G^* is *finite-in-expectation* if $l_y(h, s)$ has finite expectation for every extended history (h, s) . The strategy profile y is *finite* if there is $L_0 \in \mathbb{N}$ such that $l_y(h, s) < L_0$, for every extended history (h, s) .

An important feature of our work is the implementation of communication by pre-play correlation and short cheap-talk phases. We thus examine extensions of the game where after the first stage, players only make public announcements.

Definition 8 An almost-pre-play cheap-talk ε -equilibrium of G is an ε -equilibrium y of the game extended by cheap-talk such that at all cheap-talk phases, except at the first one, each player sends a single public message. That is : (1) the length of each cheap-talk phase is 1: $\forall h \neq \emptyset, l_y(h, s) \leq 1$, and (2) all private message are null: $\forall k > 1, n > 0, i, j \in I, s_{k,n}^{i,j} = \diamond$.

An almost-pre-play correlated ε -equilibrium of G is an ε -equilibrium of the game extended by a pre-play correlation device and by cheap-talk such that only public messages are sent: (1) $\forall h \neq \emptyset, l_y(h, s) \leq 1$, and (2) $\forall k > 1, n > 0, i, j \in I, s_{k,n}^{i,j} = \diamond$.

2.4 The main results

Our first result shows that in games with public information, a communication equilibrium payoff, or equivalently, an extensive form correlated equilibrium payoff, is an almost pre-play correlated equilibrium payoff. That is, the device may act only before the beginning of the game, provided that players can make cheap-talk public announcements throughout the game.

Theorem 9 *Let G be a game with public information, and $g \in \mathbf{R}^I$ a communication equilibrium payoff. Then g is an almost-pre-play correlated equilibrium payoff.*

The formal proof in Section 3. The intuition is as follows. The device draws all recommendations for all possible histories. Each recommendation is then encrypted using an encoding key. Player i is told the encrypted recommendations for himself, while the encoding keys (one key for each recommendation) is told to another player j . At the relevant stage, player j announces the encoding key so as to allow player i to learn the recommendation. To prevent player j announcing a false value of the key, the device authenticates the key in such a way that player i is able to tell whether the key is genuine or forged. This is done using the authentication schemes of Rabin and Ben-Or (1989), called check vectors therein.

Our second result shows that with more than two players, the mediator can be fully dispensed with. We show that, if there are at least three players, any communication equilibrium payoff is an almost-pre-play cheap-talk equilibrium payoff and a finite-in-expectation cheap-talk equilibrium payoff. Finally, if there are at least four players, it can be obtained as a finite cheap-talk equilibrium payoff.

Theorem 10 *Let G be a game with public information with three or more players, and $g \in \mathbf{R}^I$ a communication equilibrium payoff. Then,*

- (1) *g is a finite-in-expectation cheap-talk equilibrium payoff. Moreover, if there are four or more players, then g is a finite cheap-talk equilibrium payoff.*
- (2) *g is an almost-pre-play cheap-talk equilibrium payoff.*

The proof is given in Section 4. The main idea is to use the secure multiparty computation protocols of Rabin and Ben-Or (1989), and of Ben-Or et al. (1988). These protocols allow the players to replace the mediator by cheap-talk. When there are three players we adapt the protocol of Rabin and Ben-Or, which is finite-in-expectation, and when there are four or more players we adapt the protocol of Ben-Or et al., which is finite.

Remark 11 A game with public information G is finite if there exists $N_0 \in \mathbf{N}$ such that the length of each history is at most N_0 . Our proofs actually show that if G is finite, then any communication equilibrium payoff can be implemented by cheap-talk procedures which have both properties: finite-in-expectation (or finite if there are at least four players) and almost-pre-play.

Special kinds of games with public information are stochastic games. Vieille (2000a, 2000b) proved that any two-player undiscounted stochastic game (with a finite number of states) admits an equilibrium payoff (without any communication). It is an open problem whether

this is true for stochastic games with more than two players. Solan and Vieille (2002) proved that any n -player stochastic game (with a finite number of states) admits an extensive form correlated equilibrium. Our results give the following corollary:

Corollary 12 *Every undiscounted n -player stochastic game (with a finite number of states) admits a finite cheap-talk equilibrium payoff, an almost-pre-play cheap-talk equilibrium payoff, and an almost-pre-play correlated equilibrium payoff.*

Most existing literature of cheap-talk implementation deals only with finite games and with implementation of normal-form correlated equilibria, see, e.g., Forges (1990), Barany (1992), Ben-Porath (1998, 2003), Gerardi (2004), Abraham et al. (2006, 2008), and Heller (2010a). The main contribution of the present paper is the cheap-talk implementation of communication equilibria of extensive games with public information (finite and infinite).⁴

3 Proof of Theorem 9

Let G be a game with public information. From the revelation principle, we may assume without loss of generality that: (1) the device observes the history of the game and recommends actions to the players, and (2) players obediently play the recommended actions and do not send any messages. Let us fix a canonical communication device D such that the obedient profile is an ε -equilibrium of the extended game, and let g be the corresponding payoff. The canonical communication ε -equilibrium is given by a transition probability $\mu(\cdot)$ from extended histories to recommended actions. For any pair (h, s) , where h is a non-terminal history of the game and s is a history of recommendations, $\mu(h, s)$ is a probability distribution over $A(h)$.

We begin by proving the second part of Proposition 5. We define first an autonomous device D^* (which does not observe the actual history) equivalent to D . We denote by H_K the set of histories of length K ($H_0 = \{\emptyset\}$).

Step A. In a pre-play phase, D^ does the following:*

- $s_1(\emptyset) \in A(\emptyset)$ is drawn from $\mu(\emptyset)$.
- For all $a_1 \in H_1$, $s_2(a_1) \in A(a_1)$ is drawn from $\mu(s_1(\emptyset), a_1)$.
- For all $h_2 = (a_1, a_2) \in H_2$, $s_3(h_2) \in A(h_2)$ is drawn from $\mu(s_1(\emptyset), a_1, s_2(a_1), a_2)$.

⁴ Observe that Ben-Porath (1998, 2003) and Gerardi (2004) present an implementation as a sequential equilibrium (of the extended cheap-talk game), while we present only an implementation as a Nash equilibrium.

- By induction, for all $h_K = (a_1, \dots, a_K) \in H_K$, $s_{K+1}(h_K) \in A(h_K)$ is drawn from $\mu(s_1(\emptyset), a_1, \dots, s_K(h_{K-1}), a_K)$.

The construction implicitly stops when a terminal history is reached.

Step B. At the beginning of each stage $K + 1$, the device informs player i of $\{s_{K+1}^i(h_K) : h_K \in H_K\}$.

The obedient strategy of player i consists of playing $s_{K+1}^i(h_K)$ at stage $K + 1$ if the history h_K occurred and $i \in P(h_K)$. The device D^* and the obedient strategies form an extensive form correlated ε -equilibrium of the game which is equivalent to D . Indeed, at each stage $K + 1$, after history h_K , player i has the same information about other players' recommendations as under D . Moreover, player i expects all other players to obediently play the recommendations associated to h_K . He has thus the same incentives to play obediently as under D . This proves Proposition 5.

To construct an almost pre-play correlated equilibrium, we need to modify the device D^* , so that the modified device sends messages only before the start of the game. First, minimax punishments are needed in case a deviation is detected. For each player i and history h , define

$$v_h^i = \inf_{\sigma^{-i}} \sup_{y^i} \gamma_h^i(y^i, \sigma^{-i}),$$

where γ_h^i denotes the payoff of player i in the continuation game that follows h , and where the inf runs over correlated distributions of strategies of players $-i$. For each player i , and each history h , let $\sigma_*^{-i}(i, h)$ be such a distribution that achieves the infimum up to ε . The device draws $y_*^{-i}(i, h)$ according to $\sigma_*^{-i}(i, h)$ for each i, h . Each player j is informed of

$$\{y_*^j(i, h) : h \in H, i \neq j\}.$$

We describe now how the recommended actions are processed. The modified device D^{**} performs Step A as above. For each history h , fix a prime number p_h such that $p_h > |A^i(h)|$ for each i and $p_h > \frac{1}{\varepsilon}$. In the sequel, $A^i(h)$ is treated as a subset of \mathbf{Z}_{p_h} , the finite field of integers modulo p_h . We also fix for each player i a player $j(i) \neq i$. For each history h and each player $i \in P(h)$, the device does the following:

- The device draws three random variables $(x_h^i, \alpha_h^i, \beta_h^i)$ independently and uniformly distributed in \mathbf{Z}_p .
- Player i is informed of $(y_h^i := x_h^i + s^i(h), \alpha_h^i, \beta_h^i)$.
- Player $j(i)$ is informed of $(x_h^i, u_h^i := \alpha_h^i x_h^i + \beta_h^i)$.
- Each player $m \neq i, j(i)$ is informed of (α_h^i, β_h^i) .

The random draws are all done at the pre-play stage and independently across players and histories. The device then sends all these random messages to the respective players. Let us now describe the strategies of the players. After each history h (where $P(h) \neq \emptyset$):

- All the players in $\{j(i) : i \in P(h)\}$ simultaneously broadcast the pairs $(\hat{x}_h^i, \hat{u}_h^i)$. On the equilibrium path, each player broadcasts the signals he received from the device. That is, $(\hat{x}_h^i, \hat{u}_h^i) = (x_h^i, u_h^i)$.
- For each player $i \in P(h)$, all players $m \neq i$ check whether $\hat{u}_h^i = \alpha^i \hat{x}_h^i + \beta_h^i$ (they test each $j(i)$).
- If all players $j(i)$ pass their test, each player i plays $\hat{s}^i(h) := y_h^i - \hat{x}_h^i$. Then the procedure is repeated for the next history $\hat{h} = (h, (\hat{s}^i(h))_{i \in I})$.
- If a single player j does not pass a test, he is minimaxed for the remainder of the game, i.e other players play $y_*^{-j}(j, h)$.
- If several players do not pass their test, the players play arbitrarily until the end of the game.

We have thus defined a pre-play correlation device and strategies in the game with cheap-talk one-shot public announcements. Note that the induced payoff is g . Indeed, if all players use these strategies, then $\hat{s}^i(h) = s^i(h)$, thus the same actions are played under D and D^{**} . Let us now check that we have defined an 2ε -equilibrium.

Observe that player i does not get any information about the recommendations from the announcements of the device. Since α_h^i, β_h^i are uniformly distributed, independently of the rest of the game, they convey no meaningful information. Since x_h^i is uniformly distributed, independently of $s^i(h)$, so is $y_h^i = x_h^i + s^i(h)$. This latter quantity thus conveys no information about $s^i(h)$. First assume that, after each history h and for each player $i \in P(h)$, player $j(i)$ announces the true pair (x_h^i, u_h^i) . Then, each player $j(i)$ passes the test, and player i learns the value of $s^i(h) = y_h^i - x_h^i$. We have thus replicated the information structure of D^* where player i gets to learn his recommended actions for stage K at the beginning of stage K . Player i has thus the very same incentives to play the recommended actions as under D^* .

Second, let us check that for each history h and each player $i \in P(h)$, no player $j(i)$ can profitably misreport the pair (x_h^i, u_h^i) . The key point is that j does not know (α_h^i, β_h^i) and the probability of guessing this pair correctly, knowing that $u_h^i = \alpha_h^i x_h^i + \beta_h^i$, is $1/p_h$. Let $(\hat{x}_h^i, \hat{u}_h^i)$ be the pair announced by player $j(i)$. The probability of passing the test is

$$\Pr(\hat{u}_h^i = \alpha_h^i \hat{x}_h^i + \beta_h^i \mid u_h^i = \alpha_h^i x_h^i + \beta_h^i).$$

If $\hat{x}_h^i = x_h^i$ then $\hat{u}_h^i = u_h^i$ is necessary to pass the test and thus a misreport is almost surely

detected. If $\hat{x}_h^i \neq x_h^i$, the test succeeds only if $\alpha_h^i = (u_h^i - \hat{u}_h^i)(x_h^i - \hat{x}_h^i)^{-1}$ which has probability $1/p_h$. The probability to pass the test with a false report is thus at most $1/p_h$.

Such a deviation of player $j(i)$ is detected by all players with high probability. It yields player $j(i)$ an expected payoff no greater than $(1 - 1/p_h)v_h^{j(i)} + 1/p_h \leq g^{j(i)} + 2\varepsilon$. \square

Remark 13 The punishment is only needed when a deviation from the public announcements is detected. Deviations from the recommendations are already taken care of by the device, as D^{**} inherits most incentive properties of D .

Remark 14 The above proof uses the authentication schemes of Rabin and Ben-Or (1989) to process the recommended actions. Alternatively, it is possible to use a different, and in some aspects simpler, device D^{**} , which is schematically sketched as follows.⁵ Suppose that after some history h , player i has K available actions - (a_1, \dots, a_K) , and that the recommended action of the device D^* is a_k . Device D^{**} randomly and uniformly chooses K different numbers in the set $\{1, \dots, L\}$ (for large enough L), and in the pre-play phase it sends: 1) the sequence (a_1, \dots, a_K) to player i ; 2) the number a_k to player $j(i)$; and 3) a permutation of the sequence (a_1, \dots, a_K) in a random order to all other players. At the mid-play talk phase after history h , player $j(i)$ broadcasts the number a_k that he received. Assuming that player $j(i)$ followed the protocol, then player i knows his recommended action, while all other players only know that player $j(i)$ broadcasted a valid recommended action. If player $j(i)$ “lies” (broadcasts any other number), then the deviation is detected by all other players with high probability.

4 Proof of Theorem 10

The main building block of our cheap-talk implementations is the secure multiparty computation protocols of Ben-Or, Goldwasser and Wigderson (1988, henceforth BGW) and Rabin and Ben-Or (1989, henceforth RB). These protocols have been used for cheap-talk implementation of normal-form correlated equilibria in Abraham et al. (2006, 2008) and Heller (2010a). In this section we show how to adapt these protocols to implement canonical ε -communication equilibria. Specifically we use RB’s protocol when there are three players, and BGW’s protocol when there are four or more protocols.

The first subsection describes the main properties of the protocols, and in the following subsections we apply these protocols to prove the two points of Theorem 10.

⁵ We use Rabin and Ben-Or’s authentication scheme in order to make the proof of Theorem 9 more similar to the proof of Theorem 10, which extensively uses schemes of Rabin and Ben-Or (1989).

4.1 Secure Multiparty Computations

The main tool for proving Theorem 10 are the protocols of BGW and RB. The setting is as follows. Each player i knows a secret *input* $x^i \in \mathbf{Z}_p$. The aim is to jointly compute n polynomials $(f^i(x^1, \dots, x^{|I|}))_{i \in I}$, the *outputs*, in such a way that player i learns his own output $f^i(x^1, \dots, x^{|I|})$ without getting any information on the inputs and outputs of the other players. With the help of a mediator, this is very simple. Each player privately reveals his input to the mediator, the mediator computes the outputs and privately reveals $f^i(x^1, \dots, x^{|I|})$ to player i . The aim of secure multiparty computation is to construct a protocol whereby players send messages to each other and which replicate the computation by the mediator. That is, at the end of the protocol, each player i learns $f^i(x^1, \dots, x^{|I|})$, and the conditional distribution of $(x^j)_{j \in I, j \neq i}$ and $(f^j(x^1, \dots, x^{|I|}))_{j \in I}$ given the messages that player i sent and received (and his input x^i) is the same as the conditional distribution given only x^i .

The protocols of BGW and RB deal with $|I|$ players, out of which up to t players ($t < |I|/2$) may jointly deviate from the protocol. We assume $t = 1$ and $|I| \geq 3$, i.e., only unilateral deviations are possible and there are at least 3 players. The reader is referred to BGW and RB for a complete definitions of the protocols. Let us now recall the properties of these protocols that are useful to us.

Both protocols share the following *secrecy property*: a unilateral deviation does not allow the deviator to acquire any information about the inputs or the outputs of the other players.⁶

We now describe *reliability* properties of these protocols, namely the *correction property* of BGW's protocol (with four or more players), and the weaker *monitoring property* of RB's protocol. The concern is that outcomes should not be affected too much by unilateral deviations.

First, assume that there are at least four players. A strategy of player i is *obedient* if player i sends the messages recommended by the protocol. Denote $m^i(x^i)$ the obedient strategy of player i when his input is x^i . The protocol of BGW has the following *correction property*. If player i deviates and uses a strategy that is not obedient during the multiparty computation (including sending invalid messages), then his deviation is corrected in the following sense. The computation of outputs continues as if player i played according to $m^i(x^i)$, for some $x^i \in \mathbf{Z}_p$. After the protocol a *monitoring subphase* is executed: all non-deviating players

⁶ Though, when a player deviates, non-deviating players may acquire information. If a player receives an invalid message, he requires the sender to broadcast the message, and he continues the computation with respect to the broadcasted message. Thus, other non-deviating players acquire some information about inputs or outputs.

broadcast the messages sent and received during the multiparty computation (a deviator may send arbitrary messages). Since there is a strict majority of obedient players, all players agree on the values of all inputs and outputs.

Second, assume there are three players. Say that the protocol has the monitoring property if it is followed by a monitoring subphase such that (1) and (2) below are satisfied:

(1) *If only player i deviates during the multiparty computation, then all non-deviating players commonly agree that player i deviated.*

(2) *If no player deviates during the multiparty computation, then at the end of the monitoring subphase, all non-deviating players commonly agree: (i) that no deviation occurred, and (ii) on the values of the inputs and outputs of all the players.*

RB constructs a protocol that has the monitoring property with high probability. That is, for each $\delta > 0$ there exists a protocol such that for every unilateral deviation, the requirements (1) and (2) hold with probability at least $1 - \delta$. Further, if no player deviates, then (2) occurs with probability 1.

4.2 Finite Cheap-Talk Implementation

In this subsection we prove the first point of Theorem 10.

PROOF. We fix a canonical communication device D such that the obedient profile is an ε -equilibrium of the extended game. Let $\mu(\cdot)$ be the corresponding transition probability from extended histories to recommended actions, and g be the corresponding payoff. Let us construct a finite-in-expectation cheap-talk 3ε -equilibrium z that induces a payoff g_ε in an ε -neighborhood of g .

After histories $h \in H$ where $P(h) = \emptyset$, no communication is executed (players send null messages).

For each extended history (h, s) where $h \in H$ is a history of length K such that $P(h) \neq \emptyset$ and s is a history of recommendations, we construct a cheap-talk phase from which each active player $i \in P(h)$ obtains a recommended action $a^i \in A^i(h)$. The cheap-talk phase comprises three subphases: (1) monitoring of the previous stage, (2) choosing a profile by multiparty computation, (3) random monitoring (subphase (3) is needed only when there are three players). We describe how each of subphase is executed. In the following, we set $\delta(h) = \varepsilon^2/2^{K+1}$.

(1) Monitoring of the previous stage. *Each player publicly broadcasts the messages that he sent and received during the last computation subphase of the previous cheap-talk phase.*

Note that, due to the correction property (or the monitoring property when there are three players), all non-deviating players commonly agree (with probability at least $1 - \delta(h)$ when there are three players) on the profile of recommended actions that were induced in the previous cheap-talk phase, even if one of the players deviates in this subphase. As a consequence, after the extended history (h, s) , all non-deviating players agree on the value of s with probability at least $1 - \varepsilon^2$ (with probability 1 if there are at least four players).

(2) Choosing a profile. *If there is no coalition of at least $|I| - 1$ players that agree on the value of s , then players play arbitrarily in the remainder of the game.*

Otherwise, they perform a multiparty computation protocol that draws $(a^i)_{i \in P(h)}$ from the distribution $\mu(h, s)$ and informs player i of a^i only.

Note that the former case occurs with probability at most ε and only if there are three players and one of them is a deviator. In the latter case, the multiparty computation protocol is as follows.

Let $p = p(h)$ such that $p > 1/\delta(h)$, and $p > |A^i(h)|$ for each $i \in P(h)$. We assume that all action sets $A^i(h)$ are subsets of \mathbf{Z}_p and let $M(h) = \mathbf{Z}_p \cup \{\diamond\}$ be the set of messages.⁷ Let $(f^i(\cdot))_{i \in I}$ be a vector of polynomials over \mathbf{Z}_p , such that the distribution of $(f^i(x))_{i \in I}$ approximates $\mu(h, s)$ when x is uniformly distributed. Formally, the polynomials satisfy the following conditions:

- If x is uniformly distributed in \mathbf{Z}_p , then for all $(a^i)_{i \in P(h)} \in \prod_{i \in P(h)} A^i(h)$,

$$\left| \Pr \left((f^i(x))_{i \in P(h)} = (a^i)_{i \in P(h)} \right) - \mu(h, s) \left((a^i)_{i \in P(h)} \right) \right| < \delta(h).$$

- For each non-active player $i \notin P(h)$, $f^i(x) = 1$ for all $x \in \mathbf{Z}_p$.
- If $a^i \in \mathbf{Z}_p \setminus A^i(h)$ for some active player $i \in P(h)$, then

$$\Pr \left((f^i(x))_{i \in P(h)} = (a^i)_{i \in P(h)} \right) = 0.$$

Let each player i choose a uniformly distributed secret input $x^i \in \mathbf{Z}_p$ and let $x = x^1 + \dots + x^{|I|}$. As soon as at least one player i chooses x^i uniformly, then x is uniformly distributed, regardless of the way the other secrets $(x^j)_{j \in I}$ are chosen. The players then use the multiparty computation of BGW and RB for computing $(f^i(x))_{i \in P(h)}$. At the end of this subphase,

⁷ Since actions sets are finite, we can map actions one-to-one to integers in $\{1, \dots, p\}$.

each player i obtains the value of his output $f^i(x)$, which is interpreted as the protocol's recommended action for player i : if $f^i(x) = a^i$, then player i should play a^i . If some player i does not receive a valid recommended action, he chooses his action arbitrarily.⁸

If a player receives an invalid message during the computation subphase (for example, receiving a null message instead of a number in \mathbf{Z}_p), then he asks the sender to publicly broadcast the message. If the broadcasted message is invalid as well, then all non-deviating players commonly know the identity of the deviator and they minimax him for the rest of the game.

When there are four or more players, the correction property of BGW's protocol guarantees that unilateral deviations are corrected by the other players: a recommended action profile is generated according to the desired distribution, and each player correctly receives his recommended action. When there are three players we add a random monitoring subphase.⁹

(3) Random monitoring. *The players decide, according to a joint lottery, whether to perform a monitoring subphase or not. In the former case, each player broadcasts all messages he sent and received in the last computation subphase. In the latter case, nothing is revealed and the cheap-talk phase ends (every player sends null messages), and each player plays his recommended action.*

The joint lottery is conducted as follows: each player i simultaneously broadcasts a uniformly distributed random number $y^i \in \mathbf{Z}_p$. The players perform the monitoring phase if $y^1 + \dots + y^{|I|} < \varepsilon p$, which occurs with probability approximately $1 - \varepsilon$. If some player i does not broadcast a valid number, we set $y^i = 0$. The sum of the y^i 's is uniformly distributed as soon as at least one player i chooses y^i uniformly.

The monitoring property of RB's protocol guarantees that when the monitoring subphase is executed (regardless of any unilateral deviation during this subphase), with probability at least $1 - \delta(h)$, all non-deviating players correctly agree on the identity of any single deviator in the computation subphase. Assuming that they all agree that player i deviated, then all the other players minimax player i for the rest of the game: they use cheap-talk communication to implement a correlated profile that minimizes player i 's payoff.¹⁰ If no deviation was detected in the monitoring subphase, the players choose a new profile using a new computa-

⁸ This may occur only if one player deviates during the computation, and if there are exactly three players.

⁹ This subphase is an adaptation of the random monitoring presented in Ben-Porath (1998), see also Heller (2010a).

¹⁰ Recall that the non-deviating players have private communication channels that are secure from the eyes of the deviator (see Definition 6). They can use these channels to coordinate the correlated profile that minimizes the deviator's payoff.

tion subphase.

This completes the description of the cheap-talk extension and of the strategies. Now we prove that we have defined a 3ε -equilibrium that induces a payoff ε -close to g .

First, observe that if all players follow the strategies z , the distributions of actions are close to the one given by μ and thus the payoff is in an ε -neighborhood of g . Note also that by construction, z is finite (if all players follow the protocol) when there are four or more players: after each history, the players execute the finite protocol of BGW once. When there are three players, z is finite-in-expectation: each subphase is finite (due to the finiteness of RB's protocol), and the expected number of repetitions of these subphases (which are determined by the results of the joint lotteries in the random monitoring subphase) is $1/\varepsilon$, so it is finite as well.

Second, we discuss unilateral deviations. There are five types of deviations from the protocol: (1) deviation while monitoring the previous phase, (2) deviation in the computation subphase, (3) deviation in the random monitoring subphase, (4) deviation at the playing stage, (5) giving information to other players. We show that none of these deviations (nor a combination of them) is profitable to the deviator.

- (1) **Deviating at “monitoring of the previous stage” subphases.** In these subphases, players are supposed to broadcast the messages they received and sent in the previous computation subphase. Following the extended history (h, s) , player i might deviate and send different messages in this subphase. However, the monitoring/correction properties guarantee that the non-deviating players commonly agree on the value of s with probability at least $1 - \delta(h)$ (with probability 1 when there are four players or more). Thus, at all stages of the game, regardless of unilateral deviations at these subphases, all non-deviating players know the correct recommended profiles in previous stages of the game, with probability at least

$$1 - \sum_{K=1}^{\infty} \delta((a_k)_{k < K}) \leq 1 - \sum_{K=1}^{\infty} \varepsilon^2 / 2^{K+1} = 1 - \varepsilon^2.$$

Thus, with probability at least $1 - \varepsilon^2$, a deviation at these subphases is not profitable. With probability ε^2 , the deviation may not be detected, and the deviator may gain at most 1 (payoffs are between 0 and 1). Therefore, the total expected gain from these deviations is at most ε^2 .

- (2) **Deviating at computation subphases.**

- **Public deviations** - During the computation subphase, player i may broadcast an invalid message, e.g. by sending a null message instead of a number in \mathbf{Z}_p . In this case, all other players detect the deviation and minmax player i . Being minmaxed may increase player i 's payoff by at most ε relative to g , and thus by at most 2ε relative to the payoff induced by z .
- **Private deviations** - Consider first the three player case. A player may send an incorrect message, while the recipient of the message does not know that the message is incorrect. This may yield a profit of at most 1 if, at the random monitoring step, the result of the joint lottery is such that the players do not execute the monitoring subphase. However, the random monitoring is executed with probability at least $1 - \varepsilon$, and the monitoring property of the protocol guarantees that the identity of the deviator is revealed to all non-deviating players with probability at least $1 - \varepsilon^2$. In this latter case, the other players minmax the deviator for the rest of the game, and he may increase his payoff by at most 2ε . The expected gain from such a deviation is thus at most 3ε .

When there are at least four players, the correction property implies that there are no such undetected deviations: the recipient can always know whether a message is incorrect (that is, not induced by one of the protocol's obedient strategies), ask the recipient to broadcast it, and continue the computation with the broadcasted message (if it is also invalid, it is treated as a public deviation).

- (3) **Deviating in the random monitoring subphase.** We only need to consider the three player case here. Player i may deviate at the joint lottery step, but such a deviation does not change the distribution of the lottery's result. He may also deviate in the random monitoring subphase itself. The monitoring property ensure that with probability at least $1 - \delta(h)$, unilateral deviations at this stage do not affect players' assessments on deviations in the last computation subphase, and thus do not affect player i 's payoff. Thus, player i gains by deviating in a random monitoring subphase, with probability at most $\delta(h)$. As the expected number of random monitorings at each stage is $1/\varepsilon$, player i gets an expected gain of at most $\varepsilon\delta(h) = \varepsilon/2^{K+1}$, by deviating at all random monitoring subphases after a history of length K . Thus, deviating in all the random monitoring subphases throughout the game may increase the deviator's payoff by at most ε .
- (4) **Deviating in the playing stage.** Player i may play an action different from the recommended. The monitoring/correction properties imply that the other players will know the profile of past recommendations with probability at least $1 - \delta(h)$. Together

with the fact that following the device's recommendations is an ε -equilibrium of the extended game $G(D)$, this implies that player i may increase his payoff by at most $2\varepsilon + \varepsilon^2$ by the deviation.

- (5) **Giving information to other players:** Player i may deviate by sending another player (say player j) some information acquired during the computation phase, thereby allowing player j to obtain information about the recommended action profile, and have a profitable deviation (which may be also profitable to player i). Since only unilateral deviations are possible, player i should expect player j to conform with the strategies and thus disregard the extra information. When player i deviates, we are off equilibrium (we have not required any perfection properties) and thus we assume that no other player j deviates afterwards.

From this discussion, we conclude that no unilateral deviation may increase the payoff of the deviator by more than 3ε . \square

4.3 Almost-Pre-Play Cheap-Talk Implementation

In this subsection we prove the second point of Theorem 10.

PROOF. We show how to adapt the construction of the previous section to yield an almost-pre-play cheap-talk 3ε -equilibrium z' that induces a payoff in an ε -neighborhood of g . We use the same notation as in the previous proof complemented by the following: Given a history $h \in H$ of length K , let $S(h)$ be the set of histories of recommendations which are consistent with h . That is, $(s_k^i)_{k=1, \dots, K, i \in P(h_{|k})} \in S(h)$ if and only if $\forall k = 1, \dots, K, s_k^i \in A^i(h_{|k})$. For each history $h \in H$ and $s \in S(h)$, let $L(h, s) \in \mathbf{N}$ be a large enough integer such that if $L(h, s)$ many profiles are sampled according to $\mu(h, s)$, then with probability at least $1 - \delta(h) = 1 - \varepsilon/2^{K+1}$, the empirical distribution of the sample $\mu_L(h, s)$ is $\delta(h)$ -close to $\mu(h, s)$: $\forall a \in A(h), |\mu_L(h, s)(a) - \mu(h, s)(a)| < \delta(h)$.

We describe now a long pre-play cheap-talk phase, and a short public mid-play cheap-talk phases.

Pre-play communication. During the first cheap-talk phase, the players perform multiparty computation many times, to choose a large number of recommended action profiles for each possible history and each sequence of past recommended profiles. Specifically, for each extended history (h, s) where $h \in H$ and $s \in S(h)$, players execute $L(h, s)$ many times the following sub-phases: choosing a profile by multiparty computation and random monitoring. (As before, the random monitoring is executed only when there are exactly three players.

A single execution ends when the result of the joint lottery is such that the players do not conduct the random monitoring.)

At the end of this phase, the players have jointly computed $L(h, s)$ many profiles for each pair (h, s) , and each player knows only his part of each profile. If a deviation is detected in any such execution, the non-deviating players minmax the deviator. Otherwise, the players execute a single-stage public mid-play communication protocol for choosing the new action profile.

Mid-play communication. For each history $h \in H$, the players execute the following subphases: (1) For each sequence of past recommended profiles $s \in S(h)$,¹¹ the players perform a joint lottery for choosing a uniformly distributed random number $j(h, s)$ between 1 and $L(h, s)$. The recommended profile is then the $j(h, s)$ -th profile among the $L(h, s)$ profiles constructed in the pre-play cheap-talk phase. (2) The players execute a “monitoring of the previous stage” subphase, in which they simultaneously broadcast the messages of the computation of recommended action profile of the previous stage.¹² Due to the correction/monitoring property, at the end of this phase, all players know the chosen recommended profile of the previous cheap-talk phase with probability at least $1 - \delta(h)$. Thus with high probability, they commonly know s , and each player plays his $j(h, s)$ -th recommended action for (h, s) .

The same arguments as in the previous subsection imply that z' is an almost-pre-play cheap-talk 3ε -equilibrium that induces a payoff ε -close to g . □

Note that it is also possible to have an almost-pre-play cheap-talk implementation by an alternative construction,¹³ where a single recommendation profile is constructed for each extended history (instead of $L(h)$ profiles), and the players use an authentication scheme as in Section 3. Pre-play communication in this alternative construction is shorter, because players have to construct a smaller number of recommendation profiles (while the additional communication that is required to construct the authentication schemes is relatively short).

¹¹ The players commonly know all the recommended profiles except the last one.

¹² As the computation subphase is finite and bounded, players can simultaneously broadcast all these messages using a large enough finite alphabet.

¹³ We have chosen not to use this alternative construction, due to the relative complexity of its formal presentation.

5 Concluding Remarks

- (1) **Resistance to coalitional deviations:** Abraham et al. (2006, 2008) and Heller (2010a) discuss how to use Rabin and Ben-Or (1989) and Ben-Or et al. (1988)'s protocols for implementing normal-form correlated equilibria of finite games in ways that are resistant to coalitional deviations. Specifically, Heller defines a k -strong equilibrium, as a profile that is resistant to joint deviations of up to k players, and shows how to implement any k -strong normal-form correlated equilibrium as a k -strong Nash equilibrium of the extended cheap-talk game, assuming that the deviating coalition is a minority: $k < |I|/2$. The cheap-talk equilibria presented in this paper can be adapted to allow the implementation of canonical communication equilibria in games with public information in a way that is resistant to coalitional deviations of minorities.
- (2) **General action sets:** Throughout the paper we assumed that at each stage of the game each player has a finite set of actions. We now shortly discuss the extensions of our results to the case where the set of actions is a compact subset of a separable metric space. Theorem 9 can be extended to this setup. Without loss of generality, the recommended action of each player i can be represented as a sequence of zeros and ones. Each such "bit" can be encoded using the scheme described in Section 3 (where the players simultaneously send an infinite number of messages at each mid-play cheap-talk phase).

Theorem 10 can be extended only under strong continuity assumptions on the whole structure of the game tree. With such assumptions, the action profile of the players at each stage can be approximated by a finite set, and the distributed computation schemes described in Section 4 can be used. In the general case, the distributed computation schemes, which relies on operations on a finite field, cannot be used when the action sets are infinite, and we do not know whether all communication equilibrium payoffs can be obtained by unmediated cheap-talk procedures.

Acknowledgements

The research of Heller and Solan was supported by the Israel Science Foundation (grant number 212/09).

References

- [1] Abraham I., Dolev D., Gonen R., Halpern J., 2006. Distributed computing meets game theory: robust mechanisms for rational secret sharing and multiparty computation. *Proc. 25 ACM Symp. Principles of Distributed Computing*, 53-62.
- [2] Abraham I., Dolev D., Gonen R., 2008. Lower bounds on implementing robust and resilient mediators. *TCC*, 2008.
- [3] Aumann R.J., 1974. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1:67–95.
- [4] Blackwell D., Ferguson T.S., 1968. The big-match. *Annals of Mathematical Statistics* 39(1), 159-163.
- [5] Barany I., 1992. Fair distribution protocols or how the players replace fortune. *Mathematics of Operations Research*. 17, 329-340.
- [6] Ben-Or M., Goldwasser S., Wigderson A., 1988. Completeness theorems for non-cryptographic fault-tolerant distributed computation (extended abstract). *Proc. 20 STOC ACM*, 1-10.
- [7] Ben-Porath E., 1998. Communication without mediation: expanding the set of equilibrium outcomes by “cheap” pre-play procedures. *Journal of Economic Theory* 80, 108-122.
- [8] Ben-Porath E., 2003. Cheap talk in games with incomplete information. *Journal of Economic Theory* 108, 45-71.
- [9] Crawford, V., Sobel, J., 1982. Strategic information transmission. *Econometrica* 50, 579-594.
- [10] Dubey P., Kaneko M., 1984. Information patterns and Nash equilibria in extensive games: I. *Mathematical Social Sciences* 8, 111-139.
- [11] Farrell J., Rabin M., 1996. Cheap talk. *Journal of Economic Perspectives* 10, 103-118.
- [12] Forges F., 1986. An approach to communication equilibria, *Econometrica* 54, 1375-1385.
- [13] Forges F., 1990. Universal mechanisms. *Econometrica* 58, 1341-1364.
- [14] Gerardi D., 2004. Unmediated communication in games with complete and incomplete information. *Journal of Economic Theory* 114, 104-131.
- [15] Heller Y., 2010a. Minority-proof cheap-talk protocol. *Games and Economic Behavior*, doi:10.1016/j.geb.2009.11.004, in press.
- [16] Heller Y., 2010b. Sequential Correlated Equilibria in Stopping Games. mimeo. www.tau.ac.il/~helleryu/correlated-stopping.pdf

- [17] Myerson R.B., 1986. Multistage games with communication. *Econometrica* 54, 323-358.
- [18] Osborne M. J., Rubenstein A., 1994. A Course in Game Theory. The MIT Press.
- [19] Rabin T., Ben-Or M., 1989. Verifiable secret sharing and multiparty protocols with honest majority (Extended Abstract). *ACM Symp. Theory of Computing*, 73-85.
- [20] Renault J., Tomala T., 2004. Communication equilibrium payoffs in repeated games with imperfect monitoring. *Games and Economic Behavior* 49(2), 313-344.
- [21] Solan E., 2001. Characterization of correlated equilibria in stochastic games. *International Journal of Game Theory* 30, 259-277.
- [22] Solan E., Vieille N., 2002. Correlated equilibrium in stochastic games. *Games Economic Behavior*, 362-399.
- [23] Urbano A., Vila J.E, 2002. Computational complexity and communication: coordination in two-player games. *Econometrica* 70 (5), 1893-1927.
- [24] Vieille N., 2000a. Two-player stochastic games I: A reduction. *Israel Journal of Mathematics* 119, 55-91.
- [25] Vieille N., 2000b. Two-player stochastic games II: The case of recursive games. *Israel Journal of Mathematics* 119, 93-126.