

Cheap Talk in Games with Incomplete Information¹

Elchanan Ben-Porath²

January 2002

¹Running Title: Cheap Talk in Bayesian Games.

²The department of Economics and the Center for Rationality, the Hebrew University of Jerusalem, and the Eitan Berglas School of Economics, Tel-Aviv University. Correspondence: Elchanan Ben-Porath, the Center for Rationality, Feldman building, the Hebrew University, Givat-Ram, Jerusalem, Israel. e-mail: benporat@math.huji.ac.il Fax: 972-2-6513681.

Abstract

The paper studies Bayesian games which are extended by adding pre-play communication. Let Γ be a Bayesian game with full support and with three or more players. The main result is that if players can send private messages to each other and make public announcements then every communication equilibrium outcome, q , that is rational (i.e., involves probabilities that are rational numbers) can be implemented in a sequential equilibrium of a cheap talk extension of Γ , provided that the following condition is satisfied: There exists a Bayesian Nash equilibrium s in Γ such that for each type t_i of each player i the expected payoff of t_i in q is larger than the expected payoff of t_i in s .

Journal of Economic Literature classification number: C7.

Key words: communication, Bayesian games, sequential equilibrium.

1 Introduction

In a classical paper Crawford and Sobel [7] examined the implications of cheap talk in a game in which one player, the sender, has private information and another player, the receiver, has to choose an action. Crawford and Sobel show that cheap talk extends the set of outcomes that can be implemented in equilibrium by allowing for partial revelation of information. This paper addresses a general question that their paper suggests and that is: Given a general game Γ with incomplete information, what is the set of outcomes that can be generated in an extension of Γ in which players can talk before they act?

One motivation for this enquiry is a situation in which players are interested in implementing an outcome q which is not an equilibrium outcome in the original game Γ . For example q can be an outcome which Pareto-dominates every (or some) equilibria in Γ . If there exists a pre-play communication procedure which implements q in an equilibrium of the extended game, then players would want to follow it.

We now define more precisely the model and the results. The starting point is a Bayesian game Γ . We let $t_i, t_i \in T_i$, denote the private information of player i and A_i her set of actions. The game Γ is played as follows: First nature moves and chooses a profile of “types” $t = (t_1, \dots, t_n)$ according to some probability distribution $\mu \in \Delta(T)$. Each player i learns her type t_i and then in the second stage chooses an action in A_i . The choices of the actions are simultaneous. We assume that A_i and T_i are finite. An outcome q is a function which associates with each profile of types t a probability distribution on the set of profiles of actions A , $A \equiv \prod_{i=1}^n A_i$.

A cheap talk extension of Γ , $\hat{\Gamma}$, is a finite game in the extensive form in which after players learn their types they communicate. The communication can take any finite number of stages where in each stage players send simultaneously private or public messages. At the last stage of $\hat{\Gamma}$ the players choose simultaneously actions in Γ . This formulation allows, of course, for the case where at a given stage only one player, or a subset of players, sends messages (the messages of the rest are vacuous). A profile of strategies in $\hat{\Gamma}$, τ , induces an outcome q in the original game Γ . We say that q can be implemented by cheap talk if there exists a cheap talk extension $\hat{\Gamma}$ of Γ and a sequential equilibrium τ in $\hat{\Gamma}$ such that τ induces

q . Our question is what is the set of outcomes that can be implemented by cheap talk?

To describe our results we need to refer to the notion of a communication equilibrium outcome, C.E.O. A C.E.O. is an outcome that can be generated in a Nash equilibrium of an extension of Γ , Γ^d , that involves a mediator. A mediator is a device which receives private messages from the players, performs a private lottery and then sends back private messages. A C.E.O. is an extension of the notion of a correlated equilibrium distribution to games with incomplete information.¹ In particular, in a complete information game a C.E.O. is a correlated equilibrium distribution and vice-versa. It follows from the revelation principle, Myerson [15] that every C.E.O. can be generated in an equilibrium in which each player reports to the mediator her true type and the mediator sends back a message which suggests an action that the player follows. Now it is well known and easy to show that the set of outcomes that can be implemented by cheap talk is contained in the set of C.E.O.'s.² We will explain this claim in the description of the model in section 2. (In fact in section 2 we will define a C.E.O. as an outcome which is immune to a class of deviations which cannot be detected in any cheap talk extension of Γ and then show that this definition is equivalent to the definition that we have given here).

Our main result, theorem 1, is that a C.E.O. q in a game Γ can be implemented by cheap talk when the following conditions are satisfied:

- (1) The number of players in Γ is three or more.
- (2) The game Γ is a game with full support, that is, every profile of types has a positive probability.
- (3) q is rational, that is, for every profile of types t $q(t)$ involves probabilities that are rational numbers.
- (4) There exists a Bayesian Nash equilibrium s such that the expected payoff of each type t_i of each player i in s is smaller than the expected payoff of t_i in q . We call this condition the "Nash domination" condition.

Consider first condition (1). It is well known that when there are two players then in

¹The notion of a correlated equilibrium is due to Aumann [2].

²This claim is not as trivial as it might seem because the cheap talk can be conducted in many stages while the mediator operates in one stage.

games with complete information the only outcomes that can be implemented by cheap talk are convex combinations of Nash equilibria. Thus, even in the case of complete information it is possible to implement only a particular subset of the set of C.E.O.'s. (Recall that in complete information games a C.E.O. is a correlated equilibrium distribution).

Condition 2 ensures that every C.E.O. q can be implemented in a **sequential equilibrium** of the mediation game.

Condition (3) is a technical condition which has to do with joint lotteries.

Condition (4), the "Nash domination" condition, reflects the need for a credible punishment in case a player deviates from the prescriptions of the protocol of communication. The Nash equilibrium s is played when it becomes public knowledge that some player has deviated. Since the payoff in s is lower than in q the prospect of playing s keeps players in line.³ Now condition (4) can be weakened in several ways. First, it turns out that while the equilibrium does not reveal the identity of the deviator it does reveal a pair (i, j) of players from whom the deviation came. Thus, it is enough to have for every pair of players (i, j) that are revealed in case of a deviation an equilibrium $s^{i,j}$ such that each type of player i and player j prefers q upon s . In general it might be possible to do even better than that, that is, it might be possible for the players different from i and j to correlate their actions and in this manner implement a more severe punishment on i and on j . This leads us to the first extension of the main result, theorem 2, which establishes that when there are at least five players then the "Nash domination" condition can be substituted by the following (weaker) condition; For every pair of players (i, j) there exists a C.E.O. q^{ij} such that the payoff of each type of players i and j in q^{ij} is lower than its payoff in q . The C.E.O. q^{ij} is implemented if it turns out that someone from the pair (i, j) , deviated. The reason for the condition that there are at least five players in the game, (as opposed to three in our main result), is that in the implementation of q^{ij} players i and j cannot play "active" roles because there is no punishment that can be inflicted on them if one of them deviates from the protocol which generates q^{ij} .

³ Of course, the equilibrium should have the property that given a deviation by some player i there is a high probability that the fact that someone has deviated will be publicly revealed.

Our second extension, theorem 3, establishes that if public revelation of the record of communication is possible, that is, if it is possible at a predetermined point of time to publicly and credibly expose the messages that each player sent, then the Nash domination condition can be further weakened. We obtain that when the number of players is at least four, then if for every player i there exists a C.E.O. q^i such that the payoff of such type of i in q^i is lower than its payoff in q , then q can be implemented. The point here is that public exposure of the past record reveals the player who has deviated. Therefore the implementation of q can be supported by punishments that depend on the identity of the deviator.

Barany [5] examined the implications of cheap talk in games with complete information (i.e. games in the normal form). In games with complete information the only issue is the possibility of correlating actions by cheap talk, while in games with incomplete information there is in addition an issue of coordination on the basis of private information. Barany obtains the following result: Let G be a game in which the number of players is at least four and let q be a rational correlated equilibrium distribution in G . There exists a cheap-talk extension \hat{G} of G , which involves public verification of the record of communication, and a Nash equilibrium $\hat{\tau}$ in \hat{G} , such that $\hat{\tau}$ generates the distribution q .

Forges [9] shows that in a game in which the number of players is at least four, every C.E.O. is a correlated equilibrium outcome of a cheap talk extension of Γ . Combining Forge's result with the construction of Barany gives a procedure of cheap talk which can implement in a Nash equilibrium all rational C.E.O.'s. The result of Forges relates to a situation where players receive messages from a correlation device **before** they know their types and then after they learn their types communicate and take actions in Γ . Recently, Gerardi [12] has shown that it is possible to obtain Forge's result by performing the correlation at the interim stage, that is, after players learn their types. So like in the current paper all the communication takes place in the interim stage.

Comparing the results of Barany, Forges and Gerardi's extension of Forge's result, FG, to theorem 1 in the current paper we note that the former have the advantage of providing a protocol which can implement **any** rational C.E.O. while the protocol in the current paper implements only C.E.O.'s which dominate some Nash equilibrium. On the other hand, in

the former papers the protocols that are used suffer from three drawbacks:

(1) The number of players is at least four (as opposed to three in the current paper).

(2) It is assumed that the record of the communication can be publicly verified. That is, it is possible to check the content of each message that was sent.

(3) The equilibria are not sequential.

It is worthwhile to elaborate on these differences a bit. The protocol which Barany constructs and which is also used in Forges and FG deters players from deviating by using the following structure: Each message is sent by two players. If one of the players deviates and sends a message which is different from the one she should have sent then the receiver gets two different messages and she realizes that a deviation has occurred. In this case the receiver announces publicly that there is a deviation. All the messages are opened (here the assumption of the possibility of public verification is used), the identity of the deviator is revealed and the deviator is minmaxed by the other players. There are two non-sequential elements in this equilibrium. First, the profile of actions which minmaxes a player i is typically not a Nash equilibrium. (This also explains why Barany and FG do not need to assume punishment payoffs as we do). Second, it might not be in the interest of player j to announce a deviation when she receives two inconsistent messages (even if the punishment of the deviator will be with a profile of actions that is a Nash equilibrium). Finally, we note that a protocol in which each message is sent by two players requires at least four players (otherwise, one player would have known all the messages and that would make correlation by the other players impossible).

In a recent and independent work, Gerardi [10] has shown that if the number of players is greater or equal to five, then every C.E.O. can be implemented in a sequential equilibrium of a cheap talk extension of Γ . Gerardi builds a cheap talk protocol which mimics a correlation device and then uses Forges' protocol to generate any given C.E.O. The key idea in Gerardi's construction is that every message is sent by three players. If there are differences among the three senders the receiver uses a majority rule to determine the "true" message. In this manner a deviation by a single player cannot affect the outcome of the protocol. By contrast the construction in the current paper is based on a procedure of random monitoring which reveals with a high probability a deviation from a protocol

of communication. A revelation of a deviation is followed by a play of the punishment equilibrium, s . More specifically, the equilibria which we construct in the proofs of our results are composed of two ingredients:

1. A protocol of communication P which generates the given C.E.O. q .
2. A procedure of random monitoring which reveals with a high probability a deviation from the prescriptions of P and in this way ensures the execution of P in equilibrium.

The main novelty in our construction is the procedure of random monitoring which deters players from deviating from the prescriptions of P .

The research on the effects of communication in games is by now quite extensive. The following list mentions only the work that is closely related to ours and two papers which deal with applications. The notion of a communication equilibrium was first studied by Forges [8] and Myerson [15]. Aumann and Hart [4] use the notion of a bi-martingale⁴ to characterize the set of equilibrium payoffs in a cheap talk extension of a two-player game in which only one player has private information. Amitai [1] has provided a characterization of the payoff set in a two-player game with a general structure of incomplete information. As we have pointed out in a two-player game the set of payoffs that can be implemented by cheap-talk falls short of the full set of C.E.O.'s. However, recently Urbano and Villa [18] have shown that if players cannot perform computations which in the computer science literature are considered to be practically impossible, then in a two-player game every C.E.O. can be implemented in a Nash equilibrium of a cheap talk game. Lehrer and Sorin [14] approach the problem of implementation from a different direction. They show that the mediation device can be restricted without affecting its power. Specifically they show that a device which receives private messages but is deterministic and can only send public messages can still implement every C.E.O. In Ben-Porath [6] I used a procedure of repeated random verification to show that in complete information games a large class of correlated equilibrium distributions can be implemented in a sequential equilibrium of a cheap talk extension of the original game. However, the procedure that was used there would not have worked here because it would have involved revelation of private information. Furthermore, the procedure in the previous

⁴The notion of a bimartingale was first introduced in the study of repeated games with lack of information on one side (Hart [13], Aumann and Hart [3]).

paper involves public verification of the past record. Finally, I mention two papers which examine the effects of cheap talk in specific games of economic interest. Matthews and Postlewaite [16] characterize the set of equilibrium outcomes with cheap talk in a two-person double auction. They showed that this set consists of all the allocation rules which satisfy the following two conditions: 1. An ex-post individual rationality constraint which says that each type of each player upon learning the transaction which the rule specifies would still want to carry it out. 2. An adapted incentive constraint which says that each type will not benefit from misrepresenting itself even when it has the option of rejecting ex-post a transaction which the rule specifies. Furthermore, they show that every such allocation rule can be implemented in a simple cheap talk extension in which there is only one round of talk. Okuno-Fujiwara et. al [17] analyze Bayesian games in which there is one round of pre-play talk where players are restricted to exogenously specified certifiable statements, i.e. statements which do not stand in contradiction to the truth. The paper provides sufficient conditions for full revelation of private information and applies the analysis to different examples of Cournot and Bertrand duopolies.

The paper is organized as follows: The model is described in section 2. In section 3 we present a simple example which demonstrates the main ideas in our construction. In section 4 the main result, theorem 1, is proved. Section 5 discusses the extensions, theorems 2 & 3. We also show that if players can use a continuum number of messages, then any C.E.O. which satisfies the Nash domination condition and which is a convex combination of rational C.E.O.'s can be implemented.

2 The Model and the Main Result

Let $\Gamma, \Gamma = \langle N, T_i, A_i, u_i, \mu, i \in N \rangle$ be a Bayesian game where $N = \{1, \dots, n\}$ is the set of players, T_i is the set of possible types of player i , and A_i is the set of actions for player i . Let $T = \prod_{i=1}^n T_i$ and let $A = \prod_{i=1}^n A_i$; $u_i : T \times A \rightarrow R$ is a V.N.M. utility function for player i . So for $\sigma \in \Delta(A)$, $u_i(\sigma, t)$ is the utility of player i when the profile of types is t and the probability distribution on the profiles of actions is σ . We assume w.l.o.g. that $u_i(a, t) \geq 0$ for every $i \in N$, $a \in A$, and $t \in T$. Finally, μ is a probability vector on T , i.e. $\mu \in \Delta(T)$.

The sets T and A are assumed to be finite. We say that Γ is a game with full support if $\mu(t) > 0 \forall t \in T$.

The interpretation of this model is that $t_i \in T_i$ represents the private information of player i . Player i of type t_i assigns a probability of $\mu(t_{-i} | t_i)$ to the event that the profile of types of the other players is t_{-i} . A strategy s_i for player i is a function $s_i : T_i \rightarrow \Delta(A_i)$. A profile of strategies $s = (s_1, \dots, s_n)$ is a Bayesian Nash equilibrium if for every player i and type $t_i \in T_i$, $s_i(t_i)$ maximizes the expected utility of player i given her information t_i .

Formally,

$$s_i(t_i) \in \arg \max_{\sigma_i \in \Delta(A_i)} \sum_{t_{-i} \in T_{-i}} \mu(t_{-i} | t_i) u_i(\sigma_i, s_{-i}(t_{-i}), (t_i, t_{-i})).$$

Consider now some cheap talk extension of this game $\hat{\Gamma}$. The game $\hat{\Gamma}$ is a finite game in the extensive form in which players communicate after they learn their types. The communication takes some finite number of stages⁵ where at each stage players send simultaneously private or public messages. At the last stage of $\hat{\Gamma}$ the players choose simultaneously actions in the original game Γ . We note that this general formulation allows, of course, for the case where at a given stage only one player, or a subset of players, sends messages (the messages of the rest are vacuous). A profile of strategies $\tau = (\tau_1, \dots, \tau_n)$ in $\hat{\Gamma}$ induces an outcome function $q^\tau : T \rightarrow \Delta(A)$ where $q^\tau(t)$ is the probability distribution on A that is generated by τ in the state t . We say that an outcome function q can be implemented by cheap talk if there exists a cheap talk extension $\hat{\Gamma}$ of Γ and a sequential equilibrium τ in $\hat{\Gamma}$ which induces q i.e., if $q = q^\tau$. Now it is easy to see that type t_i of player i can deviate and change the outcome function q^τ without being noticed in two ways; first, by playing as though she is some other type t'_i and second by changing her actions according to some function $\delta_i : A_i \rightarrow A_i$. Therefore, if τ is an equilibrium such a deviation is not beneficial for player i of type t_i . This leads us to the following definition. Let $q : T \rightarrow \Delta(A)$ be an outcome function. We let $q(t)[a]$ denote the probability that q assigns to a in the state t , and let $u_i[q | t_i]$ denote

⁵The number of stages depends on the game Γ .

the expected payoff of type t_i w.r.t. q . That is,

$$u_i [q | t_i] \equiv \sum_{t_{-i} \in T_{-i}} \mu [t_{-i} | t_i] \sum_{a \in A} q(t_i, t_{-i}) [a] u_i(a, (t_i, t_{-i})). \quad (1)$$

We say that q is a communication equilibrium outcome, henceforth C.E.O, if

$$\begin{aligned} \forall_i \in N \forall t_i, t'_i \in T_i \text{ and } \forall \delta_i, \delta_i : A_i \rightarrow A_i \\ u_i [q | t_i] \geq \sum_{t_{-i} \in T_{-i}} \mu [t_{-i} | t_i] \sum_{a \in A} q(t'_i, t_{-i}) [a] u_i(\delta_i(a_i), a_{-i}, (t_i, t_{-i})). \end{aligned} \quad (2)$$

If we set $q = q^\tau$ then the L.H.S. is the expected payoff of type t_i of player i in the equilibrium τ , while the R.H.S. is the payoff of t_i if she plays as though she is type t'_i and changes her actions according to the function δ_i . Thus it follows that if q can be implemented by cheap talk then q is a C.E.O. We say that q is rational if for every $t \in T$ $q(t)$ is a vector of probabilities that are rational numbers. The general question that motivates our inquiry is what is the set of C.E.O's that can be implemented by cheap talk. Before we state our result formally we point out that the set of C.E.O.'s is equivalent to the set of outcome functions that is generated in a Nash equilibrium of a communication game that involves a **mediator**. Specifically, suppose that the players can use a mediator which can receive a private message $m_i \in M_i$ from each player i , conduct a private lottery, and then send back to player i a private message $r_i \in R_i$. A mediation mechanism is defined by a "strategy" of the mediator, that is a function $d, d : M \rightarrow \Delta(R)$ where $M = \prod_{i=1}^n M_i$ and $R = \prod_{i=1}^n R_i$. A mediation mechanism defines a new Bayesian game Γ^d where an action for player i is a pair (m_i, δ_i) where $\delta_i : R_i \rightarrow A_i$. Now let q be a C.E.O, q can be generated in a Bayesian equilibrium of a mediation game in which $M_i = T_i, R_i = A_i$ and $d = q$. The strategy of player i is simply to report her true type to the mediator and follow the action that the mediator recommends. We will call this strategy the 'straightforward' strategy. It is easy to see that the condition which makes the profile of straightforward strategies a Nash equilibrium in the

mediation game is exactly the condition 2 which defines a C.E.O. The revelation principle (Myerson [15]) establishes that any outcome function $q : T \rightarrow \Delta(A)$ that is generated in an equilibrium of some mediation game can be generated in a straightforward equilibrium. Therefore, the set of outcome functions that can be generated in a Nash equilibrium of the mediation game is the set of C.E.O's.

The solution notion that we use is a sequential equilibrium. Gerardi [11] constructs an example of a Bayesian game Γ without full support in which there is a C.E.O. that cannot be implemented in a sequential equilibrium of the mediation game. The point is that without full support it is possible to have a situation where player i learns from the report of the mediator that some other player has deviated at the stage where messages were sent from the players to the mediator. The notion of Nash equilibrium does not put any restriction on the behavior of a player outside the equilibrium path while sequential equilibrium does. When there is full support, player i cannot learn from the report of the mediator that a deviation has occurred and therefore it is easy to see that every Nash equilibrium in the mediation game is a sequential equilibrium as well.

We can now present our main result.

Theorem 1 *Let Γ be a Bayesian game with three or more players and full support. Let q be a rational C.E.O in Γ . If there exists an outcome function q^s that is generated by some Bayesian Nash equilibrium s , $s = (s_1, \dots, s_n)$ in Γ such that:*

$$\forall i \in N, \forall t_i \in T_i \quad u_i[q^s | t_i] < u_i[q | t_i]$$

then q can be implemented by cheap talk.

We need some additional notation and definitions. The equilibrium τ involves the execution of many private lotteries. In all the lotteries a player i selects an element at random from some finite set D according to the uniform distribution. In such a case we will simply say that player i picks at random an element $d \in D$. Given a set D we let $\Pi(D)$ denote the set of permutations (i.e., bijections) on D . In the construction of τ we make use

of a joint lottery. A joint lottery is a mechanism that generates a distribution v from private lotteries so that a deviation by a single player from executing her private lottery does not change v . The equilibrium τ involves a joint lottery which generates the uniform distribution on a set $Z = \{1, \dots, \bar{z}\}$. The mechanism is very simple, one player picks an element $z \in Z$ at random and another player picks a permutation $\pi \in \Pi(Z)$ at random. Both players announce their choices simultaneously and the outcome of the process is the element $\pi(z)$. It is easy to see that $\pi(z)$ is uniformly distributed and that a deviation by one of the players will not change that.

3 An Example

In this section we present a simple example which motivates and demonstrates the main ideas in the construction of the mechanism, i.e. an extended game and a sequential equilibrium in that game, which generates a given C.E.O. q .

There are three players, and to keep things as simple as possible, only player 1 acts and this player has no private information. The formal definition is as follows:

$$\begin{aligned}
 N &= \{1, 2, 3\} \\
 |T_1| &= 1 \quad T_2 = \{-, +\} \quad T_3 = \{-, +\} \\
 \mu(t) &= \frac{1}{4} \quad \forall t \in T \\
 A_1 &= \{a_0, a_1, a^{+,-}, a^{-,+}\}
 \end{aligned}$$

1. $u_i(a_0, t) = 0 \quad \forall i \in N \quad \forall t \in T$
2. $u_i(a_1, t) = 1 \quad i = 2, 3 \quad \forall t \in T$

$$u_1(a_1, t) = \begin{cases} 1 & t = (+, +), (-, -) \\ -2 & t = (+, -), (-, +) \end{cases}$$

$$3. \quad u_1(a^{+,-}, (+, -)) = 1$$

$$u_i(a^{+,-}, t) = -2 \quad \text{if } i \neq 1 \text{ or } t \neq (+, -)$$

$$4. \quad u_1(a^{-,+}, (-, +)) = 1$$

$$u_i(a^{-,+}, t) = -2 \quad \text{if } i \neq 1 \text{ or } t \neq (-, +).$$

Thus, the action a_0 gives everyone a moderate payoff, 0, with certainty. The action a_1 always gives players 2 and 3 a high payoff 1. It gives player 1 a high payoff in states $(+, +)$ and $(-, -)$ and a low payoff, -2, in states $(-, +)$ and $(+, -)$. The actions $a^{+,-}$ and $a^{-,+}$ are symmetric. The action $a^{+,-}$ gives player 1 a high payoff in state $(+, -)$ and a low payoff in every other state. For players 2 and 3 $a^{+,-}$ is always bad, i.e. it always gives a low payoff. Similarly, $a^{-,+}$ gives player 1 a high payoff in $(-, +)$ and a low payoff in every other state. For players 2 and 3 this action is always bad.

Note that the expected payoff for player 1 is negative for every action that is different than a_0 . Therefore there is a unique Bayesian Nash equilibrium in this game in which player 1 plays a_0 . Consider now the outcome function q which is described in figure 1.

State	Outcome
$\mathbf{t}=(\mathbf{t}_2, \mathbf{t}_3)$	$\mathbf{q}(\mathbf{t})$
$(+, +)$	a_1
$(+, -)$	a_0
$(-, +)$	a_0
$(-, -)$	a_1

Figure 1

Since the expected payoff of each type of each player is $0.5 q$ pareto dominates the Nash equilibrium a_0 . We also claim that q is a C.E.O. To see that, it is convenient to think of q as the rule which determines the recommendation of a mediator to player 1 as a function of the reports of players 2 and 3 about their types. Now it is easy to verify that: 1. Given that player 1 follows the recommendation of the mediator, it is in the interest of players 2 and 3 to truthfully reveal their types. (In fact in this example players 2 and 3 are indifferent between lying and telling the truth.) 2. Given that players 2 and 3 truthfully report their types it is in the interest of player 1 to follow the recommendation of the mediator.

Note that it is important that if player 1 receives the recommendation a_0 she doesn't know the true state - she only knows that it is either $(+, -)$, or $(-, +)$. More precisely player 1 assigns a probability 0.5 to $(+, -)$ and a probability 0.5 to $(-, +)$. If player 1 would have known that the true state is, say, $(+, -)$, she would have played $a^{+,-}$. This example demonstrates a general issue in the implementation of communication equilibrium outcomes and that is the need to reveal only **part** of the information. A mechanism will work in general only if each player does not learn anything beyond the information that is revealed by the action that is recommended to the player.

Can the players implement q without a mediator? Consider a simple cheap-talk extension $\tilde{\Gamma}$ of Γ in which in the first stage players 2 and 3 make simultaneous announcements and then in the second stage player 1 picks an action in A_1 . The following claim establishes that q cannot be implemented in $\tilde{\Gamma}$.

Claim 1: *There does not exist a sequential equilibrium τ in $\tilde{\Gamma}$ which induces the outcome function q .*

The proof, which is simple, is given in the appendix.

We note that there exists an equilibrium in which players 2 and 3 truthfully reveal their types, but in this equilibrium player 1 plays $a^{+,-}$ in state $(+, -)$ and $a^{-,+}$ in state $(-, +)$. Now it turns out that in this example there exists a cheap talk mechanism which implements q and which is simple in comparison to the general mechanism in that it involves only two stages of communication. In the first stage players 2 and 3 reveal to each other simultaneously, but not to player 1, their types. In the second stage each one of them sends a private message to player 1 recommending a_1 in case the announcements in the first stage

were $(+, +)$ and $(-, -)$, and recommending a_0 if the announcements were $(+, -)$ or $(-, +)$. If the recommendations of players 2 and 3 are identical, player 1 follows them, and if the recommendations are not identical, (which is an out-of-equilibrium event), player 1 plays a_0 . It is straightforward to verify that these strategies constitute a sequential equilibrium. However, note that during the play of the extended game player 2 learns the type of player 3 and player 3 learns the type of player 2. It is easy to construct an example of this sort (player 1 acts and players 2 and 3 have private information), in which the knowledge of players 2 and 3 about the type of the other will make the implementation of some C.E.O.'s impossible. So again, the point is that the mechanism will work in general only if it does not reveal to a player any information beyond the action that should be played.

We now turn to the construction of such a mechanism. The mechanism is composed of two elements:

1. A protocol P which generates q .
2. A monitoring procedure which provides players with incentives to follow the prescriptions of P .

We let $\hat{t} = (\hat{t}_2, \hat{t}_3)$ denote the true state of nature, $\mathcal{A}_1 : T \rightarrow A_1$ is the function that associates with each state t the action that is played by player 1 in t . So in our example $\mathcal{A}_1(t) = q(t)$.

We can now define the protocol .

THE PROTOCOL

I. Player 2 picks at random a permutation $\psi \in \Pi(T_2)$ and sends $\psi(\hat{t}_2)$ to player 3.

II. Player 2 picks at random a permutation $\sigma \in \Pi(T) = \Pi(T_2 \times T_3)$. Player 2 sends σ to player 3 and sends player 1 the function:

$$\mathcal{A}_1^{\psi, \sigma}(t) \equiv \mathcal{A}_1(\psi^{-1} \circ \sigma_2^{-1}(t), \sigma_3^{-1}(t))$$

$\mathcal{A}_1^{\psi,\sigma}$ is a “key” that will be used later on. Specifically player 1 can compute her action in \hat{t} , $\mathcal{A}_1(\hat{t})$, by applying $\mathcal{A}_1^{\psi,\sigma}$ on $\sigma(\psi(\hat{t}_2), \hat{t}_3)$.

III. Player 3 picks at random a permutation $\beta \in \Pi(T)$. She sends player 1 $\beta \circ \sigma(\psi(\hat{t}_2), \hat{t}_3)$ and sends player 2 β .

Define:

$$\begin{aligned} X_0 &\equiv \{(+, -), (-, +)\} \\ X_1 &\equiv \{(+, +), (-, -)\}. \end{aligned}$$

Thus, X_0 is the set of states in which player 1 plays a_0 and X_1 is the set of states in which she plays a_1 .

The important features of P are the following:

(1) When player 1 will learn β she will be able to compute her action in the true state $\mathcal{A}_1(\hat{t})$. She will do that by applying β^{-1} on $\beta \circ \sigma(\psi(\hat{t}_2), \hat{t}_3)$ to get $\sigma(\psi(\hat{t}_2), \hat{t}_3)$ and then applying $\mathcal{A}_1^{\psi,\sigma}$ on $\sigma(\psi(\hat{t}_2), \hat{t}_3)$ to get $\mathcal{A}_1(\hat{t})$.

(2) Furthermore, player 1 will not learn anything beyond her action, $\mathcal{A}_1(\hat{t})$.

More precisely if $\hat{t} \in X_0$ then after learning β player 1 assigns to the state $(-, +)$ the probability $\mu((-, +)/X_0) = 0.5$ and to the state $(+, -)$, she assigns the probability $\mu((+, -)/X_0) = 0.5$. Similarly if $\hat{t} \in X_1$ the probabilities that player 1 assigns to $(+, +)$ and $(-, -)$ after learning β are 0.5 and 0.5.

(3) At the end of the protocol P (but, before player 1 learns β), each player does not know anything beyond what she knew at the beginning of P . More precisely the probability distribution of each player at the end of P about the state of nature and the action of player 1 in it is the initial distribution.

Properties 2 and 3 are intuitive. Their proof relies on the following simple observation. Let Y be a set and let $y \in Y$ be an element which is selected according to the

probability $\mu \in \Delta(Y)$. Let $\sigma \in \Pi(Y)$ be a permutation on Y that is selected according to the uniform distribution on $\Pi(Y)$. Then the value of $\sigma(y)$ does not give any information on the value of y .

A complete proof of property 2 is given in the appendix (Claim 2). The proof of property 3 is similar.

It follows from properties (1) and (2) that if players 2 and 3 could have been trusted to follow the prescriptions of the protocol and to announce β at the end then that would have constituted a mechanism which implements q .⁶

The problem is that player 2 can benefit by manipulating the protocol. To see this, suppose that $\hat{t}_2 = +$. Player 2 can induce player 1 to play a_1 both at $(+, +)$ and $(+, -)$ in the following manner: Player 2 sends to player 3 $\psi = I$ and $\sigma = I$, (where I is the identity permutation), but to player 1 she sends the function $\mathcal{A}'_1, \mathcal{A}'_1 \neq \mathcal{A}_1^{\sigma, \psi}$, which is defined as follows:

$$\mathcal{A}'_1(t) = \begin{cases} a_1 & t = (+, +), (+, -) \\ a_0 & t = (-, +), (-, -) \end{cases}$$

Since $\sigma(\psi(\hat{t}_2), \hat{t}_3) = (\hat{t}_2, \hat{t}_3)$, this manipulation will indeed cause player 1 to play a_1 in $(+, +)$ and in $(+, -)$. Note that player 1 has no way of recognizing that she is manipulated because there are permutations (ψ', σ') such that $\mathcal{A}'_1 = \mathcal{A}_1^{\psi', \sigma'}$.

It follows that there is a need for a procedure which will provide incentives for the players to follow the prescriptions of the protocol. We therefore now turn to the description of the second component of the mechanism.

We start by noting that if at the end of the protocol all the players simultaneously announce their private records, that is, each player announces all the messages that she received and sent, then:

(a) If a player i lies about her record then, given that the other players are truthfully reporting their records it will become common knowledge that **someone** has lied. This

⁶In fact if that was the case we could have simplified things by omitting the selection of β , that is, substituting step 3 with a move in which player 3 sends player 1 $\sigma(\psi(\hat{t}_2), \hat{t}_3)$.

follows because every message that is received or sent by player i is sent or received by some other player j . So if player i lies about some message m there will be a mismatch between his announcement and the announcement of the player j who sent or received m .

(b) If players truthfully report their records, then a deviation of player i , $i = 2, 3$, from the prescriptions of the protocol, which involves sending to another player j a message that is inconsistent with messages that i received from, or sent to, other players, will be publicly revealed. In particular, the deviation of player 2 which we described above will be exposed because \mathcal{A}'_1 is inconsistent with $\sigma = I$. That is, $\mathcal{A}'_1 \neq A_1^{I,\psi}$ for $\psi \in \Pi(T_2)$. There are two types of deviations which will not be revealed:

(1) Player i deviates in the way she conducts her private lottery, that is, she picks a permutation according to a probability distribution which is different from the uniform distribution.

(2) Player i of type \hat{t}_i plays as a type $t_i \neq \hat{t}_i$.⁷

A deviation of the first kind does not change the outcome of the game, (the only purpose of the private lotteries is to hide information from the other players), and it follows from the definition of q as a C.E.O. that a deviation of the second kind will not benefit player i .

Thus, it follows from (a) and (b) that simultaneous announcements of the private records is a procedure which will reveal any deviation which could increase the payoff of the player. So at first sight it might seem that the following mechanism implements q : First, players play the protocol P . When the play of P is complete each player announces her private record. The announcements are made simultaneously. If the announcements reveal a deviation then the punishing Nash equilibrium s is played. If no deviation is revealed player 1 proceeds to play the action that is suggested by the protocol. Since for each type of each player the payoff in s is lower than the payoff in q it seems that this mechanism induces players to follow the prescriptions of P . However, there is a problem with this mechanism: The public revelation of the whole record exposes the private information of the players. In our example the revelation of σ and $\mathcal{A}_1^{\psi,\sigma}$ reveals ψ . Once ψ is known \hat{t}_2 can be deduced

⁷A deviation by player 2 where she sends player 1 the function $\mathcal{A}^{\sigma,\psi}$ and sends player 3 at the first stage the message $\psi'(\hat{t}_2)$ where $\psi' \neq \psi$ is, of course, equivalent to deviating by playing as a type $t_2 \neq \hat{t}_2$.

from $\psi(\hat{t}_2)$ and \hat{t}_3 can be deduced from $\sigma(\psi(\hat{t}_2), \hat{t}_3)$. So the true state \hat{t} is revealed and we have seen that when the true state is revealed player 1 will not play a_0 in $(+, -)$, and $(-, +)$.

The main new idea in the paper is to “fix” this mechanism in the following way: The players will play the protocol P many times and they will do that for each possible profile of types. (Player 1, of course, does not yet act). When all the different plays of P are complete the players conduct a joint lottery which selects a “big” subset B of the plays. (B will be specified later). For each play of P , p , that belongs to B there is a stage in which players publicly and simultaneously announce their private records in p . If, for some p , $p \in B$, these public announcements reveal that there was a deviation from the prescriptions of the protocol in the play of p then players proceed to play the punishment equilibrium s . If, on the other hand, for every p , $p \in B$, the announcements reveal compliance with P then a play \hat{p} , $\hat{p} \notin B$, which corresponds to the true profile of types $\hat{t} = (\hat{t}_2, \hat{t}_3)$ is selected, players 2 and 3 announce β , and the action $\mathcal{A}_1(\hat{t})$ is played. (If one of the players deviates and announces $\gamma \neq \beta$ then there is mismatch in the announcements and the punishment s is played. More on that later.) This is done without the players learning \hat{t} . (The way in which this is done will be specified in the formal description that follows.) Now, the point is that when a player, say, player 2, contemplates a deviation from P in the play of some p she does not know whether p will belong to the set B or not (because the set B is determined later by a joint lottery). Since the number of plays in B is close to the total number of plays and since B is chosen according to the uniform distribution the probability that p will belong to B is high and therefore a deviation from P in the play of p will be revealed with a high probability. Since the result of such a revelation will be the play of the punishment equilibrium s player 2 will lose from deviating. We note that the public announcements of the private records for each play p , $p \in B$, do not reveal the true profile of types \hat{t} because players play the protocol P the same number of times for each possible profile of types.

We now turn to a precise description of this construction. Let W be the largest expected payoff that some type t_i of some player i can make by deviating from the prescriptions of P assuming that that deviation will be unnoticed. In our example it is easy to see that only player 2 can gain by deviating from P and that the maximal expected payoff that she can get is 1. Therefore $W = 1$.

Let $Z = \{1, \dots, \bar{z}\}$ where \bar{z} satisfies the following inequality for each type t_i of each player i .

$$\frac{\bar{z} - 1}{\bar{z}} \cdot u_i[q^s | t_i] + \frac{1}{\bar{z}} \cdot W < u_i[q | t_i]. \quad (3)$$

In our example the payoff of each type of each player in the Nash equilibrium s is zero so that $u_i[q^s | t_i] = 0$. The expected payoff of each type of each player in the C.E.O. q is 0.5 so $u_i[q | t_i] = 0.5$. Since $W = 1$, \bar{z} can be 3.

With each $z \in Z = \{1, 2, 3\}$ we associate $|T| = 4$ plays of the protocol P a play $p(z, t)$ for each profile of types t (so on the whole there are 12 plays of P). It is important that when $p(z, t)$ is played $t = (t_2, t_3)$ is not known, (that is, player i only knows t_i .)

Thus, we want the players to play P for every profile $t \in T$ but without knowing which t they are playing. This is done in the following way: Let R be some order on T that is commonly known. For example R could be the order $(+, +)$, $(+, -)$, $(-, +)$, $(-, -)$. For each $z \in Z$ each player i , $i = 2, 3$ picks at random a permutation $\sigma_i^z \in P(T_i)$, where the choices of the different σ_i^z , $z \in Z$, are independent of each other. For each $z \in Z$ the four plays $p(z, t)$, $t \in T$, are conducted in the following order. Players move along the set T according to the order R and when the element $t' = (t'_2, t'_3)$ is reached they play $p(z, (\sigma_2^z(t'_2), \sigma_3^z(t'_3)))$. (So they start by playing $p(z, (\sigma_2^z(+), \sigma_3^z(+)))$ and terminate by playing $p(z, \sigma_2^z(-), \sigma_3^z(-))$. In this manner each $p(z, t)$ is played but t_i is not revealed to the players different from i . When all the plays $p(z, t)$ are complete there is a public joint lottery in which an element \hat{z} is picked at random. Then all the plays (z, t) , $t \in T$, $z \neq \hat{z}$, are checked (there are eight such plays). That is, for each (z, t) , $z \neq \hat{z}$ the players announce publicly and simultaneously their private records in the play (z, t) . If the announcements for some (z, t) reveal a deviation then the players play the punishing Nash equilibrium s and the game terminates. If for each (z, t) , $z \neq \hat{z}$, there is no deviation in the play of the protocols then players select the protocol $p(\hat{z}, \hat{t})$. The selection of (\hat{z}, \hat{t}) is simple: Each player i , $i = 2, 3$ announces $\bar{t}_i = (\sigma_i^{\hat{z}})^{-1}(\hat{t}_i)$. It follows from the definition of the process by which the player plays the different protocols $p(z, t)$ that the protocol which corresponds to

(\hat{z}, \bar{t}) is $p(\hat{z}, \hat{t})$. Also, since the permutations $\sigma_i^{\hat{z}}$, $i = 2, 3$ were picked at random, no player learns anything about the true state \hat{t} beyond what the player knew at the beginning (i.e., her type). Once $p(\hat{z}, \hat{t})$ is selected, players 2 and 3 announce β simultaneously, and then player 1 plays the action that is suggested by $p(\hat{z}, \hat{t})$, namely $\mathcal{A}_1(\hat{t})$. If players 2 or 3 deviates and announces $\gamma \neq \beta$ then there will be a mismatch in the announcements of players 2 and 3 and in such case the punishment s will be played.

To summarize the protocol is played $|Z| \times |T|$ times. When $|Z|$ is large the probability that a deviation from the prescriptions of the protocol will be revealed is high. The reason that for every $z \in Z$ there are $|T|$ plays of the protocol, a play for each $t \in T$, is to avoid a situation where the public announcements of the private records reveal information on the true types of the players.

We have completed the description of the mechanism that implements q . To verify that the strategies of the players indeed constitute an equilibrium we note that:

(1) It follows from the inequality (3) that player i , $i = 2, 3$ will not gain by deviating from the prescriptions of the protocol in one, or more, of the plays $p(z, t)$. (The l.h.s. of (3) is an upper bound on the expected payoff of player i of type t_i if she deviates, while the r.h.s. is her expected payoff if she complies).

(2) It follows from the properties of the protocol (page 15) and from the fact that the permutations $\sigma_i^{\hat{z}}$, $i = 2, 3$, are picked at random that:

(a) After the selection of $p(\hat{z}, \hat{t})$, (which is done by the announcements of \bar{t}_i), and before the announcements of the permutations β no player has obtained any information (beyond, of course, her type), on the state of nature and the action of player 1.

(b) After the announcements of β player 1 learns her action $\mathcal{A}_1(\hat{t})$ but nothing beyond that.

It follows from (a) that it is not in the interest of player i , $i = 2, 3$, to deviate at the last stage of the game by announcing a permutation $\gamma \neq \beta$. Such a deviation will lead to the play of the punishment equilibrium s while the fact that player i has no new information implies that her expected payoff if she sticks with the equilibrium and plays β is $\mu_i[q | \hat{t}_i]$.

Next, it follows from (b) and, of course, from the fact that q is a C.E.O. that it is in the interest of player 1 to play $\mathcal{A}_1(\hat{t})$.

(3) Now, suppose that player i , $i = 2, 3$, deviates at the stage in which $p(\hat{z}, \hat{t})$ is selected by misrepresenting her type, i.e. by announcing a type $\tilde{t}_i \neq \bar{t}_i$. This will lead to the implementation of the action $\mathcal{A}_1(\sigma_i^{\hat{z}}(\tilde{t}_i), \hat{t}_j)$. It follows from (a) and from the definition of q as a C.E.O. that player i will not gain from this deviation. Finally, suppose player i deviates at the stage in which the different plays of P are carried out by playing as a type t'_i when she is supposed to play as type t_i . There are two possibilities. The deviation is exposed at the stage in which the records are publicly announced. In this case the punishment s will be played. Or, the deviation is not exposed, in which case the deviation affects the outcome of the game only if it was done in the play $p(\hat{z}, \hat{t})$, (so that the outcome is $\mathcal{A}_1(t'_i, \hat{t}_j)$ instead of $\mathcal{A}_1(\hat{t})$). In this case the deviation is equivalent to the deviation we discussed previously. Thus, in any case a deviation is not worthwhile.

4 The proof of Theorem 1

We turn now to a description of the equilibrium τ for the general case.

There are two differences between the example and the general case. First, in the general case the outcome function q assigns with each profile of types t a probability distribution on A and not necessarily a single action in A like in the example. Second, in the general case each player acts and has private information. As we will see the protocol of communication can be modified so that it can incorporate these two extensions. Thus the difference between τ and the equilibrium in the example, henceforth τ^e , is the protocol of communication. Otherwise, the two equilibria are similar. In particular, the equilibrium τ , like τ^e , is built of two components: 1. A protocol of communication P which for a given profile of types t generates the distribution $q(t) \in \Delta(A)$, and 2. A monitoring procedure which provides players with incentives to follow the prescriptions of P . Furthermore, the monitoring procedure in τ is similar to the one in τ^e .

We now turn to a definition of the protocol P . Let $q : T \rightarrow \Delta(A)$ be a rational outcome function. First we note that since q is rational, there exists a natural number L such that for every profile of types $t = (t_1, \dots, t_n)$ there exists a function $\mathcal{A}_t : Y \rightarrow A$ where $Y = \{1, \dots, L\}$, such that picking at random an element $y \in Y$ and then applying \mathcal{A}_t on y generates the

distribution $q(t)$ on A . Let $\tilde{t} = (\tilde{t}_1, \dots, \tilde{t}_n)$ be a profile of types. We will describe the play of the protocol where player i plays as though she is type \tilde{t}_i . The protocol consists of 9 steps. In steps (I)-(III) there is a mechanism which leads player n to implement a lottery on the set A that corresponds to the distribution $q(\tilde{t})$, without any player obtaining any information on the types of the other players and the profile of actions that was selected in this process. This is done as follows:

I. For every $i = 1, \dots, n - 1$ player 1 chooses at random a permutation $\Psi_i \in \Pi(T_i)$. In addition player 1 chooses at random a permutation $\alpha \in \Pi(Y)$. The choices of Ψ_i and α are independent of each other. Finally player 1 reveals to players $2, \dots, n - 1$, but not to n , the permutations Ψ_i and α .

II. Player i , $i = 1, \dots, n - 1$, sends to player n the message $\Psi_i(\tilde{t}_i)$.

III. Player n picks at random an element $\bar{y} \in Y$.

Define $\Psi, \Psi : T \rightarrow T$ by $\Psi(t) \equiv (\Psi_1(t_1), \dots, \Psi_{n-1}(t_{n-1}), t_n)$. Let $\bar{t} = \Psi(\tilde{t})$ and let $\bar{x} = (\bar{t}, \bar{y})$. The profile of actions that corresponds to \bar{x} is defined as:

$$\mathcal{A}_{\Psi^{-1}(\bar{t})} \circ \alpha(\bar{y}).$$

This is the profile of actions that will be played if the current protocol is selected as the protocol to be implemented.

Note that:

1. A random choice of \bar{y} followed by an application of $\mathcal{A}_{\Psi^{-1}(\bar{t})} \circ \alpha$ generates the distribution $q(\tilde{t})$ because $\Psi^{-1}(\bar{t}) = \tilde{t}$ and because a random choice of \bar{y} generates a uniform distribution of $\alpha(\bar{y})$ on Y .

2. No player has obtained any information about the types of the other players or the value of $\alpha(\bar{y})$. Therefore no player has obtained any information about the profile of actions $a = \mathcal{A}_{\Psi^{-1}(\bar{t})} \circ \alpha(\bar{y})$.

Let $X \equiv T \times Y$. We thus define a function $\mathcal{A} : X \rightarrow A$ as follows. Let $x = (t, y)$

$$\mathcal{A}(x) \equiv \mathcal{A}_{\Psi^{-1}(t)} \circ \alpha(y).$$

In steps (IV)-(VII) there are procedures which lead to a situation where for each player i there is some message \hat{m}_i that is held by another player so that when player i learns \hat{m}_i she can compute her action, a_i . Also, players different from i do not learn anything about a_i . Step IV deals with players $1, \dots, n-1$ and steps (V)-(VII) with player n .

IV. For each $i, i = 2, \dots, n-1$, player 1 picks at random a permutation $\sigma_i \in \Pi(X)$. The choices of the different σ_i are independent. Player 1 sends player i the function $\mathcal{A}_i \circ \sigma_i^{-1}$ and sends player n all the $\sigma_i, i = 2, \dots, n-1$. Player 2 picks at random a permutation $\sigma_1 \in \Pi(X)$, she sends σ_1 to player n and sends $\mathcal{A}_1 \cdot \sigma_1^{-1}$ to player 1. (The function $\mathcal{A}_j, j = 1, \dots, n-1$ is the projection of \mathcal{A} on A_j).

At the end of step (IV) player $i, i = 1, \dots, n-1$ can compute her action, $\mathcal{A}_i(\bar{x})$, once she learns $\sigma_i(\bar{x})$ which is held by player n . Player i will be able to compute her action by applying $\mathcal{A}_i \circ \sigma_i^{-1}$ on $\sigma_i(\bar{x})$. Also, the value $\sigma_i(\bar{x})$ does not reveal to player i any new information beyond her action in \bar{x} because the permutation σ_i was chosen at random.

In steps (V)-(VII) there is a procedure that terminates in a situation where player n can learn her action $\mathcal{A}_n(\bar{x})$ if player 1 tells her the index \bar{i} of $\mathcal{A}_n(\bar{x})$ in a sequence of actions that player n receives from player 2. This is done as follows:

V. Player n picks a permutation $\sigma_n \in \Pi(X)$ at random. Player n sends player 1 $\sigma_n(\bar{x})$ and sends player 2 σ_n^{-1} .

VI. Player 1 picks at random a permutation $\varphi_1 \in \Pi(X)$ and sends it to player 2.

VII. Let R be some order on X that is common knowledge among the players. Player 2 sends to player n a message that consists of a sequence of $|X|$ actions for player n that is defined as follows. If x_i is in place $i, i = 1, \dots, |X|$ according to the order R , then the i th element in the sequence is the action of player n in the element $\sigma_n^{-1} \circ \varphi_1(x_i)$, that is, $\mathcal{A}_n \circ \sigma_n^{-1}(\varphi_1(x_i))$.

Let \bar{i} denote the index that satisfies $\varphi_1(x_{\bar{i}}) = \sigma_n(\bar{x})$. Note that $\mathcal{A}_n(\bar{x})$ is the \bar{i} th element in the sequence that player 2 sends player n so when player n learns \bar{i} she learns her

action $\mathcal{A}_n(\bar{x})$. Player n does not learn anything beyond $\mathcal{A}_n(\bar{x})$ because the permutation φ was picked at random by player 1. (Note, however, that currently player n does not know \bar{i} . She will learn it later.)

Let $\hat{m}_i = \sigma_i(\bar{x})$ for $1, \dots, n-1$, (where $\sigma_i(\bar{x})$ was defined in step IV) and let $\hat{m}_n = \bar{i}$.

At the end of step (VII) we are in a position where for each player i there is a message \hat{m}_i , which is held by some other player, such that when player i will learn \hat{m}_i she will be able to compute her action. The messages $\hat{m}_i, i = 1, \dots, n$ will be revealed in the last stage of the communication just before players choose their actions. The problem is that the player who holds the message \hat{m}_i (which is player n if $i = 1, \dots, n-1$ and player 1 if $i = n$) can misrepresent it without the other players noticing that a deviation has occurred. To deal with this problem each message \hat{m}_i is dispersed so that it is held by two players. This is done in steps (VIII) and (IX). In step (VIII) the messages $\hat{m}_i = \sigma_i(\bar{x})$ $i = 1, \dots, n-1$ are dispersed. In step (IX) the message $\hat{m}_n = \bar{i}$ is dispersed.

VIII. For every $i = 1, \dots, n-1$ player n chooses at random a permutation $\beta_i \in \Pi(X)$. Player n sends player i $\beta_i(\sigma_i(\bar{x}))$, in addition for $i = 2, \dots, n-1$ she sends β_i to player 1 and she sends β_1 to player 2.

At the end of step (VIII) player $i, i = 1, \dots, n-1$, can compute her action if she knows β_i , (she applies β_i^{-1} on $\beta_i(\sigma_i(\bar{x}))$ to get $\sigma_i(\bar{x})$ and then applies $\mathcal{A}_i \cdot \sigma_i^{-1}$ on $\sigma_i(\bar{x})$ to get $\mathcal{A}_i(\bar{x})$). Define $m_i \equiv \beta_i$. Now m_i is held by both player n and another player. So when there are public announcements of m_i no player can misrepresent m_i without contradicting another player and in this way publicly revealing a deviation.

IX. Player 1 chooses at random a permutation $\beta_n \in \Pi(\{1, \dots, |X|\})$. Player 1 sends $\beta_n(\bar{i})$ to player n and sends β_n to player 2.

At the end of step (IX) player n can compute her action if she knows β_n . So we define $m_n \equiv \beta_n$. Now the message m_n is held by both players 1 and 2. So if and when they are called to publicly announce m_n neither one of them can deviate without contradicting the other and revealing by this contradiction publicly that there was a deviation.

We have completed the description of the protocol. We now summarize its general structure and list its important properties. These properties are similar to the properties of

the protocol in the example. Given a profile of types $t = (t_1, \dots, t_n)$ a play of P generates a distribution on profiles of messages $m = (m_1, \dots, m_n)$ such that:

1. When the messages $\{m_j\}_{j \in N}$ will be made public each player i will learn her action a_i and nothing beyond this. More precisely the probability distribution of player i on the set $A_{-i} \times T_{-i}$ after she learns $\{m_j\}_{j \in N}$ is player i 's probability distribution at the beginning of the protocol conditional on her action being a_i .
2. The play of P generates the distribution $q(t)$ on A via the correspondence between $m = (m_1, \dots, m_n)$ and $a = (a_1, \dots, a_n)$.
3. The protocol P does not include the public announcements of the messages $\{m_j\}_{j \in N}$. These announcements are made at the final stage of the communication game after the monitoring procedure. At the end of P the probability distribution of each player i on $A \times T_{-i}$ is her distribution at the beginning of P . In particular player i does not learn anything about her own action a_i . Thus, at the end of P each message m_j , $j \in N$, is held by some players but the set of messages that a player holds does not reveal to him anything.

There are two additional properties of P which have to do with the possibility to recognize deviations.

4. If each player i truthfully announces her whole record, that is, the outcomes of her private lotteries, the messages she received and sent and their timings, then it becomes common knowledge whether players have followed the prescriptions of the protocol.⁸ If a player i lies about her record then the fact that someone lied will be publicly revealed.
5. If a player i deviates at the final stage of the communication game by announcing messages that are different than the messages m_j that she holds, then it will become common knowledge that someone has deviated.

Properties (2) and (5) have been demonstrated during the description of the protocol. Specifically property (2) follows because the application of the function $\mathcal{A}(x)$ generates the distribution $q(\tilde{t})$ and property (5) follows because each message m_i is held by two players so a misrepresentation of m_i by one of the players will reveal a mismatch with the announcement

⁸There are stages in the protocol in which a player performs a private lottery. It is, of course, impossible to verify that the player performed the correct lottery but since in such cases the player is indifferent among the outcomes she has no incentive to deviate.

of the other player that holds m_i . Properties (1) and (3) are similar to properties (2) and (3) of the protocol in section 3. In particular player i can compute her action a_i given m_i (see steps (VIII) and (IX)), and she does not learn anything beyond her action a_i from the revelation of the message m_j , $j \in N$. Also, the public announcement of the messages m_j , $j \in N$, is not part of the protocol P , it comes after the monitoring procedure. So at the end of P , player i 's probability distribution on $A \times T_i$ is her probability at the beginning of P . The formal proof of these properties is similar to the proof of Claim 2 in the appendix. Finally, property (4) is similar to properties (a) and (b), (pages 16-17), in the discussion of the example in section 3.

We can now turn to the final step and that is a precise definition of the structure of τ . As we have pointed out at the beginning of the section the structure of τ given the protocol P is similar to the structure of the equilibrium in section 3.

Let W be the largest payoff that some type t_i of some player i can get from deviating from the protocol. Let $Z \equiv \{1, \dots, \bar{z}\}$ where \bar{z} satisfies inequality (3) for every type t_i of player i . That is,

$$\frac{\bar{z} - 1}{\bar{z}} \cdot u_i[q^s | t_i] + \frac{1}{\bar{z}} \cdot W < u_i[q | t_i].$$

With each $z \in Z$ we associate $|T|$ plays of the protocol, a play $p(z, t)$ for each profile of types t (so on the whole there are $K = \bar{z} |T|$ plays). This is done as follows: Let R be an order on T which is common knowledge between the players. For each $z \in Z$ each player i chooses at random a permutation $\sigma_i^z \in \Pi(T_i)$ where the choices of the different σ_i^z , $z \in Z$, are independent of each other. For each $z \in Z$ the plays $p(z, t)$, $t \in T$, are conducted in the following way: Players move along the set T according to the order R and when the profile $t' = (t'_1, \dots, t'_n)$ is reached they play $p(z, \sigma^z(t'))$, that is, player i plays as type $\sigma_i^z(t'_i)$. In this way when $p(z, t)$ is played the players different from i do not know the type t_i that is played by player i . When all the plays $p(z, t)$ have been completed there is a public joint lottery in which an element $\hat{z} \in Z$ is chosen at random. Then all the plays $p(z, t)$, $t \in T$, $z \neq \hat{z}$, are checked, that is, for each (z, t') , $z \neq \hat{z}$, $t' \in T$ the players

announce publicly and simultaneously their records in the play $p(z, \sigma^z(t'))$. If for some (z, t') the announcements reveal a deviation then the players play the Nash equilibrium s and the game terminates. If for each (z, t') , $z \neq \hat{z}$ there is no deviation then each player i announces $\bar{t}_i = (\sigma_i^{\hat{z}})^{-1}(\hat{t}_i)$ where \hat{t}_i is the true type of player i . Those announcements select the play $p(\hat{z}, \hat{t})$ and now players announce simultaneously the messages m_i that correspond to this play. If these announcements reveal that someone is lying then the Nash equilibrium s is played. If no deviation is observed then each player i plays the corresponding action a_i .

We have completed the description of the structure of the equilibrium τ given the protocol. We now show that:

(i) The strategies that were described indeed constitute an equilibrium, i.e., no player has an incentive to deviate.

(ii) The equilibrium τ generates the C.E.O. q .

The proof is, in fact, identical to the proof in section 3. For the sake of a complete and clear description we give it here as well.

Property (ii) is immediate: If there is no deviation then the profile of actions that is played, a , is the profile of actions that is generated by an application of the protocol P on the true profile of types \hat{t} . It follows from property 2 of P that a is generated according to the distribution $q(\hat{t})$.

Consider now (i). When the plays, $p(z, t)$, $z \in Z$, $t \in T$ are conducted \hat{z} is uniformly distributed. It follows from property (4) of the protocol that if some player i deviates from the prescriptions of the protocol in one or more of the plays then the probability that the existence of a deviation will be publicly revealed is at least $\frac{\bar{z} - 1}{\bar{z}}$. When a deviation is revealed the Nash equilibrium s is played. Now it follows from inequality (3) that the expected payoff from the deviation is smaller than $u_i[q | t_i]$ which is the expected payoff of player i if she doesn't deviate. (Here we are also using property 3 which implies that since player i gets no information during the plays of the protocol the assessment of player i of type t_i of her payoff if she sticks to her equilibrium strategy is indeed $u_i[q | t_i]$.) Thus, it is not worthwhile to deviate from the prescriptions of P in any one of the plays $p(z, t)$. It is also not worthwhile to deviate in the final stage by announcing messages that are different from the messages produced by the protocol $p(\hat{z}, \hat{t})$ because such a deviation will be revealed by

property 5, a play of s will follow, and since $u_i[q^s \mid \hat{t}_i] < u_i[q \mid \hat{t}_i]$ player i is worse off. Next, it follows from the definition of q as a C.E.O. and from properties 1, 2 and 3 that player i of type \hat{t}_i cannot gain from misrepresenting himself by announcing an element t'_i , which is different from $(\sigma_i^z)^{-1}(\hat{t}_i)$, and, or by playing an action that is different from the action suggested by the message m_i that is generated by P . Finally, suppose player i deviates at the stage in which the different plays of P are carried out by playing as a type t'_i when she is supposed to play as type t_i . If the deviation is exposed at the stage in which the records are publicly revealed then the punishment s will be played. If the deviation is not exposed then the only case in which it affects the outcome of the game is when the play that is implemented is $p(\hat{z}, (\hat{t}_{-i}, t'_i))$ where $\hat{t} = (\hat{t}_i, \dots, \hat{t}_n)$ is the true profile of types and $t'_i \neq \hat{t}_i$. As we have just seen it follows from the definition of q as a C.E.O. and from properties 1, 2 and 3 that player i of type \hat{t}_i does not gain from such a deviation.

We have verified that a player cannot gain by deviating. Thus, we have completed the proof that τ is an equilibrium which implements the outcome q .

5 Extensions

Theorem 1 does not provide a complete characterization of the set of outcomes that can be implemented by cheap talk. In this section we discuss three extensions of theorem 1 which make an additional step in that direction. The first two results, theorems 2 and 3, show that with additional assumptions it is possible to weaken the “Nash domination” condition in theorem 1, that is the condition that there exists a Nash equilibrium s such that the payoff of each type of each player in s is lower than her payoff in the C.E.O. to be implemented, q . The third extension is technical. We show that if players use a continuum number of messages, then every C.E.O. q which is a convex combination of rational C.E.O.’s and which satisfies the Nash domination condition can be implemented by cheap talk. Our first result, theorem 2, establishes that when there are at least five players a C.E.O. q can be implemented if for every pair of players, i, j there is a C.E.O. $q^{i,j}$ such that the payoff of each type of players i and j in $q^{i,j}$ is lower than her payoff at q . Formally:

Theorem 2 *Let Γ be a Bayesian game with at least five players and let \bar{q} be a rational C.E.O. in Γ such that for every pair of players i, j , there exists a C.E.O. $q^{i,j}$ which satisfies $\forall t_k \in T_k, k = i, j, u_k[q^{i,j} | t_k] < u_k[\bar{q} | t_k]$. Then \bar{q} can be implemented by cheap talk.*

Note that the Nash domination condition in theorem 1 is weakened in two ways. First, it is not necessary to have an equilibrium that punishes all the players. Second, the punishments do not have to be Nash equilibria, they can be C.E.O.'s.

Our second result, theorem 3, assumes the possibility of public verification of the past record. This means that at a given point of time, players can credibly reveal their records so everyone sees the actual messages that were sent at each stage. We have:

Theorem 3 *Let Γ be a Bayesian game and let \bar{q} be a rational C.E.O. in Γ such that for every player i there exists a C.E.O. q^i which satisfies:*

$$\forall t_i \in T_i \quad u_i[q^i | t_i] < u_i[\bar{q} | t_i]..$$

If there are four or more players in Γ and if public verification of the record of communication is possible then \bar{q} can be implemented by cheap talk.

Note that in theorem 3 the ‘‘Nash domination’’ condition has been further weakened in that each player may have a different punishment.

The proofs of theorems 2 and 3 are similar. It is more convenient to start with theorem 3. Since the construction of the equilibrium $\bar{\tau}$ which implements a C.E.O. \bar{q} that satisfies the conditions of the theorem is a simple extension of the construction in section 4 we will not repeat all the details but just outline the main steps.

The first observation is that public exposure of the record of a play p of the protocol P publicly reveals the identity of each player that has deviated from the prescriptions of the protocol because in contrast to the case where there is only cheap talk here everyone learns the **actual** messages that have been sent. The general structure of $\bar{\tau}$ is as follows: Let $Q = \{\bar{q}, q^1, \dots, q^n\}$. Let $P^q, q \in Q$, denote the protocol which implements q . For each $q \in Q$, the protocol P^q is played many times and for each profile of types just like in the proof of

theorem 1. Specifically for each $q \in Q$ the protocol P^q is played $\bar{z} | T |$ times where \bar{z} is large enough so that it satisfies the following condition: For every $i \in N$ and $t_i \in T_i$:

$$\frac{\bar{z} - 1}{\bar{z}} \cdot u_i[q^i | t_i] + \frac{1}{\bar{z}} \cdot W < u_i[\bar{q} | t_i]. \quad (4)$$

Thus, for every (q, z, t) , $q \in Q$, $z \in Z \equiv \{1, \dots, \bar{z}\}$, $t \in T$ there is a play $p(q, z, t)$ of the protocol P^q . When all these plays are complete there is a joint lottery in which an element $\hat{z} \in Z$ is picked at random. Then all the records in all the plays $p(q, z, t)$ $z \neq \hat{z}$, $q \in Q$, $t \in T$, are publicly revealed. If there is no deviation in all the records that are revealed then players proceed to implement the outcome of the play $p(q, \hat{z}, \hat{t})$ where \hat{t} is the true profile of types. This is done in a way which is similar to the play in the equilibrium τ . Suppose now that the public exposure of the records reveals that some players have deviated. Let i be the player who made the first deviation w.r.t. the order in which the different plays $p(q, z, t)$ were conducted. In this case players will proceed to implement the outcome of the play $p(q^i, \hat{z}, \hat{t})$, that is, they will implement the C.E.O. q^i which punishes player i . It follows from inequality (4) that player i should not deviate. Now the reader who has not been completely exhausted yet might notice that there are two difficulties that should be addressed:

1. What happens if player i deviates in the plays that correspond to his punishment q^i ? Maybe he can turn his punishment into a reward!
2. What if player i deviates at the stage in which the messages m_j , $j = 1, \dots, n$ are announced. As the equilibrium stands the identity of i is not revealed. The only thing that is revealed is that player i contradicts some other player j .

Both these difficulties can be overcome if the number of players is at least four. To see that, consider again the protocol P in section 4. Note that there are only three “active” players - 1, 2 and n . While the only thing that other players do is report their respective types (encoded) in step 1. In other words the only way in which a player j , $j \neq 1, 2, n$, can effect the outcome of the protocol is by misrepresenting her type and/or following an action which is different from the action that the protocol prescribes. It follows from the

definition of a C.E.O. that such deviations are not beneficial. Now, since there are at least four players we can assign roles to players in the protocol that implements q^i so that the “active” players are different from i . In such a protocol player i will have no incentive to deviate. We have addressed the first difficulty. The second difficulty can also be easily solved using again the fact that there are at least four players. This is done by modifying the protocol as follows: In step 8 player n sends the message β_i to **two** different players (rather than just one). Similarly in step 9 player 1 sends β_n to two different players. Now in the final stage of the game each message $m_i = \beta_i$, $i = 1, \dots, n$ is held by **three** players. Let i be a specific player and let j, k and ℓ denote the players which hold the message m_i . At the final stage of the game when player i hears the announcements of j, k and ℓ player i takes the announcement that was made by the majority as the message m_i . Since there are three players which are announcing the message a deviation by one of them will not change the fact that m_i is announced by the majority.

We have completed the proof of theorem 3. The proof of theorem 2 is similar in nature. The difference is that the identity of a deviator is not revealed. Specifically, if a player i deviates by sending player j a message that is different from the message which the protocol prescribes then, (assuming that player i will try to cover her deviation), public revelation of the private records only reveals that the deviator is either player i or player j .

Therefore we construct an equilibrium $\tilde{\tau}$ which implements \bar{q} as follows. With every C.E.O. $q^{i,j}$ we assign a protocol $P^{q^{i,j}}$ which implements $q^{i,j}$ and in which the “active” players are different from i and j . This is possible because we are assuming that there are at least five players. The structure of the equilibrium is now clear. Let $\tilde{Q} = \{q^{i,j} \mid i, j \in N\} \cup \{\bar{q}\}$. For each $q \in \tilde{Q}$ the protocol which corresponds to q , P^q is played $\tilde{z} \mid T \mid$ times where \tilde{z} satisfies the following condition for every $i, j \in N$ and $t_i \in T_i$:

$$\frac{\tilde{z} - 1}{\tilde{z}} \cdot u_i[q^{i,j} \mid t_i] + \frac{1}{\tilde{z}}W < u_i[\bar{q} \mid t_i].$$

When all the plays $p(q, z, t)$, $q \in Q$, $z \in Z = \{1, \dots, \tilde{z}\}$, $t \in T$ are complete there is a joint lottery in which an element $\hat{z} \in Z$ is picked at random. Then there are public

announcements of all the private records in all the plays $p(q, z, t)$, $z \neq \hat{z}$. If the announcements do not reveal any deviation the players proceed to implement the outcome of $p(q, \hat{z}, \hat{t})$ where \hat{t} is the true profile of types. Otherwise, let $i, j \in N$ be the first pair of players that contradict each other according to the order of the moves in which the game was played. In this case players will proceed to implement the outcome of the protocol $p(q^{i,j}, \hat{z}, \hat{t})$.

It is easy to see that the arguments which were given in the proof of theorem 2 establish that $\tilde{\tau}$ is an equilibrium.

Finally, consider theorem 1 again. We will show now that if players can use a continuum number of messages then any C.E.O. \hat{q} which is a convex combination of rational C.E.O's. and which satisfies the Nash domination condition can be implemented in cheap talk. So suppose that $\hat{q} = \sum_{\ell=1}^k \alpha_{\ell} \hat{q}_{\ell}$ where each \hat{q}_{ℓ} is a rational C.E.O, $0 \leq \alpha_{\ell} \leq 1$, $\ell = 1, \dots, k$ and $\sum_{\ell=1}^k \alpha_{\ell} = 1$. Let s be Bayesian Nash equilibrium which is dominated by \hat{q} .

First, we claim that we can assume w.l.o.g. that each \hat{q}_{ℓ} dominates s . This follows because we can choose the C.E.O.'s \hat{q}_{ℓ} to be close to \hat{q} . The formal proof of the claim is omitted.⁹ To implement the C.E.O. \hat{q} we extend the mechanism which was described in the proof of theorem 1 as follows: At the beginning of the extended game, players 1 and 2 conduct a joint public lottery which selects the C.E.O. \hat{q}_{ℓ} with probability α_{ℓ} . This is done as follows: Let $0 = x_1 < x_2, \dots, < x_{k+1} = 1$ be a partition of the segment $[0, 1]$ to intervals such that $x_{\ell+1} - x_{\ell} = \alpha_{\ell}$, $\ell = 1, \dots, k$. Player i , $i = 1, 2$, conducts a private lottery in which she chooses a number $\alpha_i \in [0, 1]$ according to the uniform distribution. Then players 1 and 2 announce publicly and simultaneously the numbers α_1, α_2 . If $\alpha_1 + \alpha_2 \pmod{1} \in [x_{\ell}, x_{\ell+1})$ then players will play the mechanism which implements \hat{q}_{ℓ} . This procedure implements \hat{q} .

⁹The proof can be obtained from the author upon request.

APPENDIX

Proof of Claim 1 in Section 3: Let $\hat{\Gamma}$ be an extended game in which in the first stage players 2 and 3 make simultaneous announcements and then in the second stage player 1 picks an action in A_1 . Let M_i , $i = 2, 3$ be the set of messages of player i . A strategy for player 1 is a function $\tau_1, \tau_1 : M \rightarrow \Delta(A_1)$, where $M \equiv M_2 \times M_3$ and a strategy for player i , $i = 2, 3$, is a function $\tau_i : T_i \rightarrow \Delta(M_i)$. We want to show that if $\tau = (\tau_1, \tau_2, \tau_3)$ is a sequential equilibrium then τ does not implement q . Assume by contradiction that it does. First, we note that for every pair of announcements $m = (m_2, m_3)$ which is made with a positive probability $\tau_1(m)$ assigns probability 1 to one action in A_1 . (Otherwise τ would not implement q). Let m be a pair of announcements such that $\tau_1(m) = a_0$. Let $\tau_i(t_i)[m_i]$ denote the probability assigned by τ_i to m_i when the type of i is t_i .

If τ implements q then:

$$\tau_2(+)[m_2] \cdot \tau_3(+)[m_3] = 0. \quad (\text{A.1})$$

Otherwise, player 1 plays a_0 with a positive probability in $(+, +)$. Similarly:

$$\tau_2(-)[m_2] \cdot \tau_3(-)[m_3] = 0. \quad (\text{A.2})$$

Since m is played with a positive probability either:

$$\tau_2(-)[m_2] \cdot \tau_3(+)[m_3] > 0 \quad (\text{A.3})$$

or

$$\tau_2(+)[m_2] \cdot \tau_3(-)[m_3] > 0. \quad (\text{A.4})$$

We claim that if τ is an equilibrium both inequalities should be satisfied. To see this, assume w.l.o.g. that $\tau_2(+)[m_2] \cdot \tau_3(-)[m_3] = 0$. In such a case the optimal action for player 1 in m is $a^{-,+}$ because m is announced only in the state $(-, +)$. This contradicts our assumption that τ is an equilibrium. It follows that inequalities (A.3) and (A.4) are satisfied but this implies that $\tau_i(+)[m_i] > 0$ and $\tau_i(-)[m_i] > 0$, $i = 2, 3$ which contradicts inequalities (A.1) and (A.2). Thus, our assumption that there exists a sequential equilibrium τ that implements q has been contradicted.

We turn now to Claim 2 (Section 3). To formulate the claim and prove it we need to refer to a space Ω where an element $w \in \Omega$ describes not only the profile of types of the players, but also the record that is generated by the play of the protocol, (where by record we mean all the messages that are exchanged). Let P_r denote the probability distribution on Ω that is generated by picking a profile of types $t \in T$ according to the distribution μ and then running the protocol, P . Let F denote the record of player 1 after learning β . F is a random variable and a value of F , f , is a triple $(\mathcal{A}'_1, \beta, \beta(t'))$ where $\mathcal{A}'_1 = \mathcal{A}_1^{\psi, \sigma}$ and $t' = \sigma(\psi(\hat{t}_2), \hat{t}_3)$. Now it is easy to see that since the permutation β is picked at random the informational content of $(\mathcal{A}'_1, \beta, \beta(t'))$ is equivalent to the information given by the pair (\mathcal{A}'_1, t') . So to simplify matters we can think of f as (\mathcal{A}'_1, t') . Let $\mathcal{A}'_1(t') = \alpha$ where $\alpha \in \{a_o, a_1\}$. We want to show that after learning β player 1 does not know anything beyond her action α . To be completely formal our claim 2 is:

Claim 2: *Let $f = (\mathcal{A}'_1, t')$ if $\mathcal{A}'_1(t') = a_o$, then $\Pr(t | f) = \mu(t | X_o)$ for every $t \in T$. Similarly, if $\mathcal{A}'_1(t') = a_1$, then $\Pr(t | f) = \mu(t | X_1)$ for every $t \in T$.*

Proof: We will prove the proposition for the case where $\mathcal{A}'_1(t') = a_o$. The proof for the case $\mathcal{A}'_1(t') = a_1$ is similar. Now it is clear from the definition of the protocol that if $t \in X_1$, then $\Pr(t | f) = 0 = \mu(t | X_o)$. So what we need to show is that for $t \in X_o$:

$$\Pr(t | f) = \mu(t | X_o) = 0.5 \tag{A.5}$$

We will show that:

$$\Pr(f \mid (+, -)) = \Pr(f \mid (-, +)) \quad (\text{A.6})$$

(A.5) follows from (A.6) by Bayes law (recall that $X_o = \{(+, -), (-, +)\}$).

We will prove (A.6) by showing that given any \mathcal{A}'_1 and t' such that $\mathcal{A}'_1(t') = a_0$ there is a one-to-one correspondence λ between pairs of permutations $(\hat{\psi}, \hat{\sigma})$ which give rise to the record $f = (\mathcal{A}'_1, t')$ in the state $(+, -)$. and pairs of permutations $(\bar{\psi}, \bar{\sigma})$ which give rise to f in the state $(-, +)$. More formally we will show that there is a one-to-one correspondence λ between pairs $(\hat{\psi}, \hat{\sigma})$ such that:

$$\mathcal{A}'_1^{\hat{\psi}, \hat{\sigma}} = \mathcal{A}'_1 \quad \text{and} \quad \hat{\sigma}(\hat{\psi}(+), -) = t'.$$

and pairs $(\bar{\psi}, \bar{\sigma})$ such that:

$$\mathcal{A}'_1^{\bar{\psi}, \bar{\sigma}} = \mathcal{A}'_1 \quad \text{and} \quad \bar{\sigma}(\bar{\psi}(-), +) = t'.$$

Since the permutations ψ and σ are picked independently according to the uniform distribution establishing the correspondence λ proves (A.6).

The correspondence λ is simple. Let \tilde{t} denote the state $\hat{\sigma}(\hat{\psi}(-), +)$:

$$\lambda(\hat{\psi}, \hat{\sigma}) \equiv (\bar{\psi}, \bar{\sigma})$$

where $\bar{\psi} \equiv \hat{\psi}$ and where $\bar{\sigma}$ is defined as follows:

$$\begin{aligned}\bar{\sigma}(\hat{\psi}(+), -) &= \tilde{t} \\ \bar{\sigma}(\hat{\psi}(-), +) &= t'.\end{aligned}$$

On the two elements of T that are different from $(\hat{\psi}(+), -)$ and $(\hat{\psi}(-), +)$, $\bar{\sigma}$ equals $\hat{\sigma}$.

References

- [1] M. Amitai, Cheap talk with incomplete information on both sides, discussion paper #90, Center for Rationality, The Hebrew University of Jerusalem, 1996.
- [2] R. Aumann, Subjectivity and correlation in randomized strategies, *J. Math. Econ.* **1** (1974), 67-96.
- [3] R. Aumann and S.Hart, Bi-Convexity and Bi-Martingales, *Israel Journal of Mathematics* vol.54 No.2 (1986), 159-180.
- [4] R. Aumann and S. Hart, "Long negotiations, cheap talk," Mimeo, Center for Rationality, The Hebrew University of Jerusalem, 1996.
- [5] I. Barany, Fair distribution protocols of how the players replace fortune, *Math. Oper. Res.* **17** (1992), 329-340.
- [6] E. Ben-Porath, Correlation without medication: Expanding the set of equilibrium outcomes by "cheap" pre-play procedures, *J. Econ. Theory* **80** (1998), 108-122.
- [7] V. Crawford and J. Sobel, Strategic information transmission, *Econometrica* **50** (1982), 579-594.
- [8] F. Forges, An approach to communication equilibria, *Econometrica* **54** (1986), 1375-1385.
- [9] F. Forges, Universal mechanisms, *Econometrica* **58** (1990), 1341-1364.
- [10] D.Gerardi, Unmediated Communication in Games with Complete and Incomplete Information, mimeo, Northwestern University, July, 2000.
- [11] D.Gerardi, Communication equilibria and sequential rationality, mimeo, Northwestern University, 2000.
- [12] D.Gerardi, "Interim Pre-Play Communication," mimeo, Northwestern University, September, 2000.

- [13] S.Hart, "Nonzero-sum two person repeated games with incomplete information," *Math. Oper. Res.* **10** (1985), 117-153.
- [14] E. Lehrer and S. Sorin, One-shot public mediated talk, *Games Econ. Behav.* **20** (1997), 131-149.
- [15] R. Myerson, Optimal coordination mechanisms in general principle-agent problems, *J. Math. Econ.* **10** (1982), 67-82.
- [16] S. Matthews and A. Postlewaite, Preplay communication in two-person sealed-bid double auctions, *J. Econ. Theory* **48** (1989), 239-263.
- [17] M. Okuno-Fujiwara, A. Postlewaite and K. Suzumura, Strategic information revelation, *Rev. Econ. Stud.* **57** (1990), 25-47.
- [18] A. Urbano and J. E. Villa, "Unmediated Talk under Incomplete Information", Mimeo 1997, The Valencian Institute for Economic Research.

ABSTRACT

The paper studies Bayesian Games which are extended by adding pre-play communication. Let Γ be a Bayesian game with full support and with three or more players. The main result is that if players can send private messages to each other and make public announcements then every communication equilibrium outcome, q , that is rational (i.e., involves probabilities which are rational numbers), can be implemented in a sequential equilibrium of a cheap talk extension of Γ , provided that the following condition is satisfied: There exists a Bayesian Nash equilibrium s in Γ such that for each type t_i of each player i the expected payoff of t_i in q is larger than the expected payoff of t_i in s .