Interim Pre-Play Communication*

Dino Gerardi

Northwestern University†

September 2000

1 Communication Equilibria without Mediation

In an important contribution, Forges (1990) shows that any communication equilibrium of a finite Bayesian game with four or more players and rational parameters[1] can be implemented without the help of a mediator. Starting from a finite Bayesian game, Forges constructs an extended game in which players exchange "cheap" messages before playing the original game. Messages are cheap in the sense that they do not affect *directly* players' payoffs. Forges demonstrates that the payoffs associated with any communication equilibrium are Nash equilibrium payoffs of the original game extended by adding a two phase universal mechanism[2] of ex-ante and interim pre-play plain conversation.

To obtain this strong result, Forges proceeds in two steps. First, she demonstrates that a communication equilibrium payoff of a game with incomplete information (with at least four players) is a correlated equilibrium payoff of an extended game in which the players, after learning their types, exchange cheap messages and then play the

*I thank Françoise Forges for her helpful comments.

†Department of Economics, 2003 Sheridan Road, Evanston, IL 60208-2600. E-mail: d-gerardi@nwu.edu.

[1] A Bayesian game has rational parameters if the payoffs and the probability of each profile of types are rational numbers.

[2] This means that the same mechanism can be used to implement *any* communication equilibrium of *any* game with at least four players and rational parameters.

original game. Note that in a communication equilibrium, messages go both from the players to the mediator (players report their types to the mediator) and from the mediator to the players (the mediator suggests an action to each player). Thus, the mediator helps the players exchange their private information, and coordinate their actions. In contrast, in a correlated equilibrium, messages go only from the mediator to the players (the mediator sends each player a strategy, i.e., a function that specifies an action for every possible type). The role of the mediator is to coordinate players' actions. The first result of Forges shows how players can transmit information with plain conversation, provided that there is a reliable mediator who coordinates their actions.

Forges then shows that players do not need an outside correlation device, in order to implement a communication equilibrium. Forges applies a result due to Barany (1992) to the extended game she constructs in the first step. Barany (1992) considers finite games with complete information and at least four players. He shows that the payoff associated with a correlated equilibrium is a Nash equilibrium payoff of an extended game obtained by adding a communication phase to the original game.

To summarize, the game proposed by Forges for implementing a communication equilibrium is as follows. In the ex-ante stage (i.e. before players learn their types), players exchange cheap messages (this is the communication phase proposed by Barany (1992)). Then players learn their types. After that, at the interim stage, players undergo a new round of communication (this is introduced by Forges). Finally, the original game is played.

The problem with this construction is that players make moves before and after learning their types. When we consider games with incomplete information, we usually assume that, first, Nature randomly selects a profile of types and informs each player of her own type, then players choose their actions. As Forges acknowledges, "in the contest of games with incomplete information, one might want pre-play conversation to be performed entirely at the interim stage (the ex ante stage is often interpreted as virtual)" (Forges (1990), p. 1350). Let us explain why Barany's communication phase cannot be performed at the interim stage. In Barany (1992), each message is sent by two players to a third one. Each receiver has the power to stop the communication phase. In that case, the messages (which are recorded) become publicly known and the deviator is identified.[3] The deviator is punished by her opponents who play a "minmax" strategy against her in the original game. This punishment cannot be performed at the interim stage, since the "minmax" strategy against a deviator depends on her type, which is unknown to her opponents. If Barany's scheme is played at the ex-ante stage, then a deviator can be easily punished at her *expected* minmax level.

However, in this note we show that it is possible to implement a communication equilibrium even if players communicate only at the interim stage.

Theorem 1 *Let G be a finite game of incomplete information, with at least four players and rational parameters. Every communication equilibrium payoff of G is a Bayesian-Nash equilibrium payoff of G extended by a (two phase) universal mechanism of (interim) pre-play plain conversation.*

Given the result of Proposition 1 of Forges (1990),[4] we only need to show that given a game G with incomplete information and four or more players, a *correlated* equilibrium payoff of G is a Bayesian-Nash equilibrium payoff of G extended by a mechanism of *interim* pre-play communication. In the next section, we show how this can be done when the correlated equilibrium is rational, i.e., it involves probabilities that are rational numbers. Given this, any communication equilibrium that can be expressed as a convex combination of rational communication equilibria can easily be implemented by using "jointly controlled lotteries" (see Forges (1990) for details).

We now provide an informal description of our construction (a formal proof is in the next section). To implement a rational correlated equilibrium of G, we combine two communication schemes. The first one is very similar to Barany's (1992) scheme. If players follow it, they select a strategy profile of G at random, according to the correlated equilibrium distribution, and each player learns only her own strategy. In this communication scheme the behavior of a player is *independent* of her type. In this way, no private information is revealed, and every player has an incentive to play the

[3]The deviator can be either a player who sends a wrong message or the receiver who stops the communication phase after receiving the same message from the two senders.

[4]This result states that every communication payoff of a game G with incomplete information and at least four players is a correlated equilibrium payoff of G extended by a mechanism of interim pre-play plain conversation.

strategy that she learns. The novelty of our construction is the second communication scheme, which is designed to prevent unilateral deviations in the first one. If a player deviates in the first communication scheme, messages are traced back, and the chosen strategy profile and the identity of the deviator become publicly known. Let f denote the deviator's strategy in the selected profile. A player other than the deviator plays the role of the mediator and selects a new strategy profile at random, according to the correlated equilibrium distribution *conditional* on the strategy f. Then she reports the corresponding strategy to every player. Note that the player who plays the role of the mediator selects strategies of G and not actions, and therefore she does not need to know her opponents' types. In equilibrium, every player who is not a deviator follows the "mediator's" recommendation. Since the strategy profile is selected according to the correlated equilibrium distribution conditional on f, the optimal action for the deviator is to choose the action specified by f for her type. Thus, unilateral deviations are not profitable. Finally, note that the behavioral strategies of the "mediator" are off the equilibrium path, and we do not put any restriction on them since our solution concept is Bayesian-Nash equilibrium.[5]

2 From Correlated Equilibria to Bayesian-Nash Equilibria

In this section we consider games of incomplete information with four or more players. We show that any rational correlated equilibrium can be implemented without a mediator, even though communication takes place entirely at the interim stage.

Let $G = \langle P_1, ..., P_n, T_1, ..., T_n, C_1, ..., C_n, p, u_1, ..., u_n \rangle$ be a finite Bayesian game. Players are $P_1, ..., P_n$, and we assume $n \geqslant 4$. C_i denotes the set of actions of P_i, and $C = \prod_{i=1}^{n} C_i$ is the set of action profiles. T_i is the set of types of P_i and p is a probability distribution over $T = \prod_{i=1}^{n} T_i$, the set of profiles of types.[6] We assume that for every player P_i and for every type t_i in T_i, the probability $p(t_i)$ is strictly

[5] The question of which correlated or communication equilibria can be implemented without a mediator using the stronger solution concept of sequential equilibrium is addressed by Ben-Porath (1998, 2000) and Gerardi (2000).

[6] For notational simplicity, we assume that beliefs in G are consistent.

positive (i.e. we delete types that occur with probability zero). We let $T_{-i} = \prod_{j \neq i} T_i$ denote the set of profiles of types of players different from P_i. Finally, the payoff function of P_i is $u_i : C \times T \to \mathbb{R}$.

A strategy for P_i is a function that assigns an action in C_i to every type in T_i. The set of strategies of P_i is A_i (a_i denotes the generic element of A_i). $A = \prod_{i=1}^{n} A_i$ is the set of strategy profiles and $\Delta(A)$ is the set of probability distributions over A. Given a probability distribution $r \in \Delta(A)$, we use $r(a_{-i}|a_i)$ to denote the conditional probability of a_{-i} given the strategy a_i.[7] The marginal probability of a_i is denoted by $r(a_i)$.

A probability distribution $r \in \Delta(A)$ is a correlated equilibrium of G if and only if:

$$\begin{aligned} &\sum_{t_{-i} \in T_{-i}} \sum_{a_{-i} \in A_{-i}} p(t_{-i}|t_i)\, r(a_{-i}|a_i)\, u_i(a_{-i}(t_{-i}), a_i(t_i), t_{-i}, t_i) \geqslant \\ &\quad \sum_{t_{-i} \in T_{-i}} \sum_{a_{-i} \in A_{-i}} p(t_{-i}|t_i)\, r(a_{-i}|a_i)\, u_i(a_{-i}(t_{-i}), c_i, t_{-i}, t_i), \\ &i = 1, ..., n, \quad \forall t_i \in T_i, \quad \forall a_i \in A_i \text{ with } r(a_i) > 0, \quad \forall c_i \in C_i. \end{aligned} \tag{1}$$

We let $u^r = (u_i^r(t_i))_{t_i \in T_i, i=1,\ldots,n}$ denote the vector of payoffs associated with correlated equilibrium r. The expected payoff of player P_i when her type is t_i and she receives the strategy a_i is denoted by $u_i^r(t_i, a_i)$ and is equal to the L.H.S. of inequality (1).

A correlated equilibrium is rational if for every strategy profile a in A, the probability $r(a)$ is a rational number.

To implement a correlated equilibrium without the help of a mediator, we extend the original game G by allowing unmediated communication between the players. Specifically, we consider extended games in which players first learn their types, then exchange cheap messages, and finally, play the original game G.

Proposition 1 *Let G be a finite Bayesian game with four or more players. Let u^r be the vector of payoffs associated with a rational correlated equilibrium r. There exists an extended game $\bar{G}(r)$ such that u^r is a Bayesian-Nash equilibrium payoff of $\bar{G}(r)$.*

Proof. Our proof relies heavily on Barany (1992). Given a rational correlated equilibrium r, we construct an extended game $\bar{G}(r)$ and a Bayesian-Nash equilibrium

[7] As usual, we use a_{-i} to denote a strategy profile of players different from P_i. A_{-i} is the set of all such profiles.

of $\bar{G}(r)$ that induces the distribution r on A.

We start by illustrating how the players generate the probability distribution r. Since r is rational, there exists a positive integer m and, for every a in A, a non-negative integer m_a such that $r(a) = \frac{m_a}{m}$. Define $X = \{1, ..., m\}$, and let $\{X_a\}_{a \in A}$ be a partition of X such that $|X_a| = m_a$ for every a in A. For $i = 1, ..., n$, let the projection $pr_i : A \to A_i$ be defined by $pr_i(a) = a_i$ if $a = (a_1, ..., a_i, ..., a_n)$. We extend each projection pr_i to X as follows:

$$pr_i(x) = a_i \quad \text{if } x \in X_a \quad \text{and} \quad pr_i(a) = a_i.$$

It is easy to verify that if an element $y \in X$ is selected at random, according to the uniform distribution, and if every player P_i chooses the strategy $pr_i(y)$, then every strategy profile a in A is selected with probability $r(a)$.

The set X and the projections $pr_1, ..., pr_n$ are common knowledge among the players.

The extended game $\bar{G}(r)$ consists of several steps. Players communicate with each other by the exchange of written messages, and they can verify publicly the past record of these messages. Suppose that in Step l, player P_i sends a message to P_j. Player P_j reads the note and then places it at a specific point which corresponds to the triple (i, j, l). As long as the communication process is not stopped, the note remains closed, so the other players do not observe its content. However, if the communication process is stopped, the note is opened in front of all players and its content is publicly revealed.

We are now ready to present the extended game $\bar{G}(r)$. For each step of $\bar{G}(r)$, we also describe equilibrium behavioral strategies (hereafter simply called equilibrium strategies). Then, we prove that unilateral deviations are not profitable, and therefore that the strategy profile that we propose is indeed a Bayesian-Nash equilibrium of $\bar{G}(r)$. Since the game is played at the interim stage, in every step we need to specify equilibrium behavior for every type of each player. However, to simplify the exposition, we adopt the following convention. If in a step, all types of a player behave in the same way, we will refer to that player without mentioning her type.[8]

[8] For example, if we say that player P_i sends message v, we mean that *all* types of P_i send the message v.

Step 0. *Random choices.*

In this step, which consists of different substeps, players "jointly select" random permutations on X and a random element of X. We will explain later how the joint selections are made. We first list the random choices that players make. P_1 and P_2 jointly select permutation α, P_2 and P_3 choose permutation β, P_3 and P_4 select permutation γ, and P_1 and P_4 choose permutation δ. Moreover, P_1 and P_3 jointly select permutation σ, and P_2 and P_4 choose an element $x \in X$. Finally, for $i = 1, ..., 4$, players in the set $\{P_1, P_2, P_3, P_4\} \setminus \{P_i\}$ choose permutation τ_i.

It is convenient to summarize Step 0 in Table 1, where we list the random choices known to every player P_i $(i = 1, ..., 4)$.

Player	Random Choices									
P_1	α			δ	σ			τ_2	τ_3	τ_4
P_2	α	β				x	τ_1		τ_3	τ_4
P_3		β	γ		σ		τ_1	τ_2		τ_4
P_4			γ	δ		x	τ_1	τ_2	τ_3	

Table 1: Random Choices of Step 0

There is a separate substep for each random choice. Every choice is made according to the uniform distribution over the underlying probability space, and each choice is made independently of all others.

We now describe how the players jointly select a random permutation, or a random element of X. Consider, for example, the substep in which P_1, P_2 and P_3 have to select the random permutation τ_4. The two players with the lowest indices (P_1 and P_2 in this case) make two announcements simultaneously. Specifically, P_i, $i = 1, 2$, announces a permutation τ_4^i on X to the players in the set $\{P_1, P_2, P_3\} \setminus \{P_i\}$. The chosen permutation will be $\tau_4 = \tau_4^1 \tau_4^2$ (note that τ_4 is common knowledge among P_1, P_2 and P_3). In equilibrium, P_i selects a permutation τ_4^i at random, according to the uniform distribution. This implies that the random permutation τ_4 also is distributed uniformly.

A similar procedure is used to make the remaining choices.[9] In every substep,

[9] To select the random element x, P_2 and P_4 simultaneously announce $x^1 \in X$ and $x^2 \in X$,

the two players who have to make an announcement choose their messages randomly, according to the uniform distribution on the underlying probability space, independently of the messages that they have already sent or received.

The state $y = \sigma(x)$ determines the strategy profile of the original game G that the players will play in the final step of $\bar{G}(r)$ (Step 4). In equilibrium, y has a uniform distribution over X, since σ and x are uniformly distributed. Note that no player knows the realization of y at the end of Step 0. However, Steps 1-3 are carefully designed to make every player P_i learn $pr_i(y)$ and nothing more.

Steps 1-3 consist of different substeps. In each substep, two players send two messages to a third player. After receiving the two messages, the receiver has two options. She has to decide whether to continue or to stop the communication process. If she decides to continue, the game proceeds to the next step. If the receiver of Substep l decides to stop the communication process, all messages are publicly revealed and the game proceeds to Substep $D.l$ (see below).

Substep 1.1. Player P_i, $i = 2, 3$, sends a permutation on X to P_1. The two senders report their messages simultaneously. Let b be a permutation on X. The meaning of message b in Substep 1.1 is "The realization of the random permutation $\beta\tau_1$ is b". To simplify the presentation we say that in Substep 1.1, P_2 and P_3 send the permutation $\beta\tau_1$ to P_1 (we adopt this terminology to describe the next steps).

Equilibrium strategies prescribe that each type of each sender reports the realization of $\beta\tau_1$ to P_1. For example, if the realizations of the random permutations β and τ_1 are $\hat{\beta}$ and $\widehat{\tau_1}$, respectively, in equilibrium P_2 and P_3 send the message $\hat{\beta}\widehat{\tau_1}$ to P_1.

The equilibrium strategy of P_1 is to stop the communication process if P_2 and P_3 send two different messages.

Substep 1.2. Players P_3 and P_4 send the permutation $\gamma\tau_1$ to P_1. In equilibrium, the two senders report the realization of $\gamma\tau_1$, and P_1 stops the communication process only if she receives two different messages.

Note that in equilibrium, P_1 learns the realization of the permutation $\beta\gamma^{-1} = \beta\tau_1(\gamma\tau_1)^{-1}$. However, she does not know τ_1, and so she does not learn the realizations of β and γ.

Substeps 1.3-1.8 are very similar to Substeps 1.1 and 1.2. Specifically, in Substep

respectively. The chosen element is $x = x^1 + x^2$, where $+$ is $\bmod m$.

1.3, P_3 and P_4 send the permutation $\gamma\tau_2$ to P_2. In Substep 1.4, P_2 receives $\delta\tau_2$ from P_1 and P_4. P_1 and P_2 send the permutation $\alpha\tau_3$ to P_3 in Substep 1.5, and P_1 and P_4 report the permutation $\delta\tau_3$ to P_3 in Substep 1.6. Finally, P_4 receives the permutation $\alpha\tau_4$ from P_1 and P_2 in Substep 1.7, and the permutation $\beta\tau_4$ from P_2 and P_3 in Substep 1.8. In equilibrium, the senders are honest and the receivers stop the communication process only if their senders disagree.

Substep 1.9. P_2 receives the permutation $\delta\sigma$ from P_1, and the permutation $\gamma\sigma$ from P_3. In equilibrium, P_1 and P_3 are honest and report the realizations of $\delta\sigma$ and $\gamma\sigma$, respectively.

Suppose that P_1 sends message b and P_3 sends message b'. We say that the two messages are *consistent* if $b\,(b')^{-1} = \delta\gamma^{-1}$. Note that P_2 knows the realization of $\delta\gamma^{-1}$ from Substeps 1.3 and 1.4. In equilibrium, she stops the communication process only if the two messages are not consistent.

Substep 1.10. This step is similar to the previous one. P_4 receives the permutation $\alpha\sigma$ from P_1, and the permutation $\beta\sigma$ from P_3. Equilibrium strategies prescribe that both senders are honest and that P_4 stops the communication process only if the messages of P_1 and P_3 are not consistent.[10]

Substep 2.1. P_1 receives $\gamma\sigma\,(x)$ from P_2 and $\beta\sigma\,(x)$ from P_4. Note that P_2 learns $\gamma\sigma$ in Substep 1.9 and x in Step 0. In equilibrium, she reports the realization of $\gamma\sigma\,(x)$ to P_1. Similarly, P_4 learns $\beta\sigma$ in Substep 1.10 and x in Step 0. P_4's equilibrium strategy is to send P_1 the realization of $\beta\sigma\,(x)$.

With a slight abuse of notation, let b and b' denote the messages of P_4 and P_2, respectively. The two messages are consistent if $b = \beta\gamma^{-1}\,(b')$. P_1 can tell whether the two messages are consistent, since she learns the permutation $\beta\gamma^{-1}$ in Substeps 1.1 and 1.2. In equilibrium, P_1 stops the communication process if the two messages are not consistent.

Substep 2.2. This step is similar to Substep 2.1. P_3 receives $\delta\sigma\,(x)$ from P_2 and $\alpha\sigma\,(x)$ from P_4.[11] In equilibrium, the two senders are honest, and P_3 stops the

[10] Let b and b' denote the messages of P_1 and P_3, respectively. The two messages are consistent if $b\,(b')^{-1} = \alpha\beta^{-1}$. Notice that P_4 learns $\alpha\beta^{-1}$ in Substeps 1.7 and 1.8.

[11] Note that P_2 learns $\delta\sigma$ in Substep 1.9 and P_4 learns $\alpha\sigma$ in Substep 1.10. Both players know x from Step 0.

communication process if the two messages are not consistent.[12]

Step 2 ends here if there are exactly four players. Otherwise Step 2 continues as follows. For each player P_k (with $k > 4$), there is step in which P_1 and P_2 report $\gamma\sigma(x)$ to P_k.[13] In equilibrium, both P_1 and P_2 are honest, and P_k stops the communication process only if the two senders disagree.

Substep 3.1. P_1 receives the mapping $pr_1\gamma^{-1}$ from P_3 and P_4. The equilibrium strategies prescribe that the two senders report the realization of $pr_1\gamma^{-1}$, and that P_1 stops the communication process if she receives two different messages. Note that P_1 already knows $\gamma\sigma(x)$ and so she learns $pr_1(y) = pr_1\gamma^{-1}\gamma\sigma(x)$.

In Substeps 3.2-3.4, players P_2, P_3 and P_4 learn $pr_2(y)$, $pr_3(y)$ and $pr_4(y)$, respectively. Specifically, in Substep 3.2, P_1 and P_4 send the function $pr_2\delta^{-1}$ to P_2. In Substep 3.3, P_3 receives $pr_3\alpha^{-1}$ from P_1 and P_2. In Substep 3.4, P_2 and P_3 report $pr_4\beta^{-1}$ to P_4. If the number of players is greater than four, then for each player P_k (with $k > 4$), there is a step in which P_k receives $pr_k\gamma^{-1}$ from P_3 and P_4 (so P_k learns $pr_k(y)$). As usual, in equilibrium all senders are honest and the receivers stop the communication process only if they receive two different messages.

If the communication process is never stopped, the game $\bar{G}(r)$ proceeds to Step 4.

Step 4. In this step, all players simultaneously choose an action (P_i selects an action in C_i) and the extended game $\bar{G}(r)$ ends. The equilibrium behavior of each player depends on her type. Consider player P_i of type t_i. In equilibrium, she plays the action $(pr_i(y))(t_i)$.[14]

To complete the description of $\bar{G}(r)$, let us assume that the receiver of Substep l stops the communication process. As mentioned before, in this case the game proceeds to Substep $D.l$. Suppose that the receiver of Substep l is P_i and that the two senders are P_j and P_k. Substep $D.l$ proceeds as follows. Let h denote the smallest positive integer different from i, j and k. Player P_h sends a message to every player. The set of messages to $P_{h'}$ is $A_{h'}$, the set of pure strategies of $P_{h'}$ in G. After P_h sends her messages, all the players simultaneously choose an action ($P_{h'}$ chooses an action

[12] As usual, let b and b' denote the messages of P_4 and P_2. The two messages are consistent if $b = \alpha\delta^{-1}(b')$. Notice that P_3 learns the permutation $\alpha\delta^{-1}$ in Substeps 1.5 and 1.6.

[13] P_1 learns $\gamma\sigma(x)$ in Substep 2.1.

[14] Note that P_i learns the function $pr_i(y)$ in Step 3.

in $C_{h'}$) and the game $\bar{G}(r)$ ends.

We now describe equilibrium strategies in Substep $D.l$. Since messages of Step 0 - Substep l are publicly verified, the chosen profile $(pr_1(y), ..., pr_n(y))$ becomes publicly known and the players who did not follow their equilibrium strategies (deviators) are identified. Let P_ι be the deviator with the smallest index (among P_i, P_j and P_k).[15] In equilibrium, P_h selects a strategy profile of G (i.e. an element of A) according to the following probability distribution:

$$q(a) = \begin{cases} r(a_{-\iota}|pr_\iota(y)) & \text{if } a = (a_{-\iota}, pr_\iota(y)) \\ 0 & \text{otherwise.} \end{cases}$$

Then P_h reports the h'th component of the selected strategy profile to player $P_{h'}$ $(h' = 1, ..., n)$. Equilibrium strategies prescribe that every type of each player who did not deviate follows P_h's recommendation.[16]

We now show that the strategy profile of $\bar{G}(r)$ that we have described constitutes a Bayesian-Nash equilibrium. We start by showing that unilateral deviations in Step 4 are not profitable. Consider P_i of type t_i, and suppose that she learns (in Step 3) that her strategy is $a_i = pr_i(y)$. Since her opponents do not condition their behavior on their types in Steps 0-3, P_i does not learn anything about her opponents' types. In other words, the conditional probability that the vector of the opponents' types is t_{-i}, given the information that P_i has in Step 4, is still $p(t_{-i}|t_i)$. Moreover, given the equilibrium strategies, the fact that P_i learns the strategy a_i implies that the probability that her opponents play the strategy profile a_{-i} is $r(a_{-i}|a_i)$ (see Barany (1992) for details). Since r is a correlated equilibrium, type t_i of P_i has no incentive to deviate from $a_i(t_i)$.

We now consider unilateral deviations in Steps 1-3. Consider P_i of type t_i, and remember that, given her opponents' strategies, she does not have any new information about their types. We need to consider two possible situations: (i) P_i does not know $pr_i(y)$; (ii) P_i knows the realization of $pr_i(y)$, which we denote by $\overline{a_i}$. We analyze the two cases separately.

[15] Note that if a receiver stops the communication process, there is at least one deviator: if the senders did not deviate the receiver is the deviator.

[16] Suppose that $P_{h'}$ is not a deviator, that she has type $t_{h'}$ and that she receives strategy $a_{h'}$ from P_h. In equilibrium, $P_{h'}$ chooses the action $a_{h'}(t_{h'})$.

Case (i).

The continuation payoff of type t_i of P_i is:

$$\sum_{t_{-i}\in T_{-i}}\sum_{a\in A}p\left(t_{-i}|t_i\right)r\left(a\right)u_i\left(a_{-i}\left(t_{-i}\right),a_i\left(t_i\right),t_{-i},t_i\right)=\sum_{\{a_i\in A_i:r(a_i)>0\}}r\left(a_i\right)u_i^r\left(t_i,a_i\right).$$

If P_i deviates, then she will be identified as the unique deviator, and the chosen strategy profile will be revealed. With probability $r\left(a_i\right)$ the strategy chosen for player P_i will be a_i and a lottery on $a_i\times A_{-i}$ will occur, according to the probability distribution $r\left(a_{-i}|a_i\right)$. Given her opponents' strategies, type t_i of P_i can obtain at most $u_i^r\left(t_i,a_i\right)$ by playing $a_i\left(t_i\right)$. Therefore, the maximum payoff that t_i can achieve with a unilateral deviation is $\sum_{\{a_i\in A_i:r(a_i)>0\}}r\left(a_i\right)u_i^r\left(t_i,a_i\right)$, which prevents t_i from deviating.

Case (ii)

If type t_i of P_i knows the strategy $\overline{a_i}$, her continuation payoff is $u_i^r\left(t_i,\overline{a_i}\right)$. If she deviates, it becomes publicly known that P_i is the deviator and that $pr_i\left(y\right)=\overline{a_i}$. Again, a lottery on $\overline{a_i}\times A_{-i}$ will occur, according to the probability distribution $r\left(a_{-i}|\overline{a_i}\right)$ and type t_i will be able to get at most $u_i^r\left(t_i,\overline{a_i}\right)$. Thus, the deviation is unprofitable.

Finally, unilateral deviations in Step 0 do not have any consequences. Given the equilibrium strategies of P_i's opponents, every random choice of Step 0 is independent of P_i's messages and uniformly distributed over the underlying probability space. Thus P_i is indifferent among all possible messages in Step 0. ■

The extended game $\bar{G}\left(r\right)$ that we have presented in the proof depends on the game G and the correlated equilibrium r that we want to implement. In particular, the set of messages that players use in the different steps of $\bar{G}\left(r\right)$ are functions of G and r. However, our mechanism of pre-play communication can easily be made universal by taking the set of messages of each step of $\bar{G}\left(r\right)$ to be equal to $\mathbb{N}$, the set of natural numbers. In this way, any rational correlated equilibrium of any game with four or more players can be implemented by using the *same* extended game. Clearly, the dependence on the parameters of the game and on the correlated equilibrium that one wishes to implement is through the equilibrium strategies of the extended game.

References

[1] Barany, I. (1992): "Fair Distribution Protocols or How the Players Replace Fortune," *Mathematics of Operations Research*, 17, 327-340.

[2] Ben-Porath, E. (1998): "Correlation without Mediation: Expanding the Set of Equilibrium Outcomes by "Cheap" Pre-play Procedures," *Journal of Economic Theory*, 80, 108-122.

[3] Ben-Porath, E. (2000): "Cheap Talk in Games with Incomplete Information," Working Paper No. 8, The Foerder Institute for Economic Research, Tel-Aviv University.

[4] Forges, F. (1990): "Universal Mechanisms," *Econometrica*, 58, 1341-1364.

[5] Gerardi, D. (2000): "Unmediated Communication in Games with Complete and Incomplete Information," Northwestern University (mimeo).