

Internal Correlation in Repeated Games¹

By E. Lehrer²

Abstract: This paper characterizes the set of all the Nash equilibrium payoffs in two player repeated games where the signal that the players get after each stage is either trivial (does not reveal any information) or standard (the signal is the pair of actions played). It turns out that if the information is not always trivial then the set of all the Nash equilibrium payoffs coincides with the set of the correlated equilibrium payoffs. In particular, any correlated equilibrium payoff of the one shot game is also a Nash equilibrium payoff of the repeated game.

For the proof we develop a scheme by which two players can generate any correlation device, using the signaling structure of the game. We present strategies with which the players internally correlate their actions without the need of an exogenous mediator.

1 Introduction

The well-known folk theorem (see [A2]) suggests that the set of all the Nash equilibrium payoffs in infinite undiscounted repeated games coincides with the set of all the feasible and individually rational payoffs (FIR). Any such payoff can be sustained by a frequency strategy, in which the relative frequency of any payoff tends to its weight in the convex combination. When one of the players deviates, all other players punish him and push his payoff down to his individually rational level. For such a strategy to be properly carried out, it is necessary that any deviation be detectable. That is, the actions should be observable. Such an information structure is commonly called standard information. However, if the information is not standard, some deviations can go unnoticed. Thus, some points in FIR may not be sustainable by equilibrium. The question of characterizing the set of (upper) Nash equilibrium payoffs in general undiscounted repeated games with nonstandard information is still unanswered. We provide here a characterization of the Nash equilibrium payoffs in a family of games with nonstandard information. It turns out that in this family of games the sets of Nash equilibrium payoffs and of correlated equilibrium payoffs are closely related.

¹ The author acknowledges J.-F. Mertens and S. Sorin and two anonymous referees of the *International Journal of Game Theory* for their comments and suggestions.

² Ehud Lehrer, Department of Managerial Economics and Decision Sciences, J. L. Kellogg Graduate School of Management and Department of Mathematics, Northwestern University, 2001 Sheridan Road, Evanston, Illinois 60208, USA.

The concept of correlated equilibrium, introduced by Aumann [A1], allows for an external mediator who provides the players with private information before the game starts. Usually, such an extension of a game (i.e., by adding an exogenous device) enlarges the set of equilibrium payoffs.

In repeated games with standard information the set of the correlated equilibrium payoffs (CE) coincides with the set of the Nash equilibrium payoffs (NE). In other words, any equilibrium payoff supported by an external correlation device can be achieved as an equilibrium payoff even without an exogenous mediator. This is not surprising, though, because CE contains NE and the latter already contains all the feasible payoffs, which are individually rational.

A game whose NE coincides with CE is called a *saturated game*. It is clear that any two player repeated game, not necessarily with standard information, for which NE coincides with FIR is a saturated game. However, in repeated games where the information is not standard NE is typically smaller than FIR, and CE is larger than NE. The question naturally arises: It is possible to find an information structure for which the respective repeated game (with any payoff matrix) is saturated and, moreover, the corresponding NE is typically smaller than FIR?

In this paper we introduce a family of repeated games, called games with $S-T$ (for standard-trivial) information. In these games the signal a player gets is either revealing the action combination played (standard) or completely concealing it (trivial). In symmetric $S-T$ games the signal is standard for one player if it is standard for the other. The signal in symmetric $S-T$ games can be conceived of also as a commonly observed signal which is either the action combination played or a null signal that does not reveal any information.

We show that symmetric $S-T$ games are saturated. Along with the characterization of CE in general repeated games with nonstandard information (see [L2]) this result provides also a characterization of NE in the case of symmetric $S-T$ games.

Notice that unless the signal is standard, a player is not explicitly informed of his own payoff. The payoff is deposited in the player's bank account to which he has no access during the game. A vast majority of the literature devoted to repeated games with incomplete information or to repeated games with imperfect monitoring deals with models in which a few or all players are not explicitly informed of their own payoffs. For a state-of-the-art exposition of repeated games with and without complete information, the reader is referred to the forthcoming book of Mertens-Sorin-Zamir [MSZ].

This paper joins the growing body of papers dealing with repeated strategic interactions with imperfect monitoring. Radner [R], and Rubinstein-Yaari [RY], have studied undiscounted repeated games with a one-sided moral hazard and two players a principal and an agent. The agent is fully informed and the principal is partially informed of the agent's moves. Fudenberg-Maskin [FM], and Abreu-Pearce-Stachetti [APS] have studied discounted games in which after each stage all of the players are informed of a common signal which depends (perhaps stochastically) on the action combination played. Fudenberg-Levine [FL] dealt with an n -player repeated game with observable payoffs. They defined a set of mutually punishable and enforceable payoffs (i.e., payoffs that are associated with joint actions from which any profitable deviation is detectable) and showed that these payoffs are equilibrium payoffs of the repeated game. Some of the techniques

employed here and in the literature referred to resemble those used by Kohlberg [K]. In [K], Kohlberg studies zero-sum repeated games with both incomplete information and general information functions for the two players.

The typical equilibrium payoff in undiscounted games with complete information has the following structure: first, it should be sustained by a combination of actions from which any profitable deviation is detectable and, second, it should be punishable. For instance, in the case of standard information, any feasible payoff is associated with a combination of joint pure actions from which no one can deviate without being detected and, if it is individually rational, it is also punishable. This is the structure of equilibrium payoffs of some games with imperfect monitoring (i.e., nonstandard information) as well. The reader is referred to [L1], where the case of observable payoffs is treated to [L4], which deals with the lower Nash equilibrium, and to [L5] in which the semi-standard case is dealt with. The strategies sustaining such payoffs usually consist of two phases. The first one is the master plan in which the players play an action combination from which any profitable deviation is detectable. In the second phase, the punishment plan, players can communicate among themselves in order to transmit information regarding previous deviations, and to punish the defector, if they find an alleged deviation.

The main purpose of this paper is to present a new type of strategy, one with a third phase – the correlation phase. In this phase the players generate by their own moves a correlation device to be used in the master phase.

In games with S – T information, payoffs sustainable by external correlation devices, are also Nash equilibrium payoffs. We prove that any external correlation can be substituted by an internal correlation which utilizes only the information structure of the game. By using it properly, the players can generate during the correlation phase any correlation matrix according to which they play in subsequent periods.

The proof relies mainly on two previous results. The first one is a characterization of CE in repeated games with imperfect monitoring (see [L2]). The second result is that limits of certain *finitely* repeated games payoffs (those associated with strategies from which any profitable deviation is detectable) are sustainable by equilibria in the *infinitely* repeated game (see [L3]). The method is to show that any correlated equilibrium payoff of the repeated games is a limit of such payoffs.

The paper is built as follows. In Section 2 we give the general model of repeated games with symmetric S – T information, and we present the notion of saturated games. The third section is devoted to the analysis of a specific example. In the fourth section we introduce the main theorem. The fifth section contains pertinent results from [L2], [L3] and [L5]. The sixth section considers the case where there is a player, all of whose deviations are detectable.

In the seventh section we demonstrate by an elaborate example how the information structure can be used to create any exogenous correlation. In other words, we demonstrate how any correlation matrix can be generated during the game. This section can be also regarded as a part of the preplay communication literature. Barany [B] and Forges [F] demonstrated the possibility of designing a correlating procedure which ends up with a correlation distribution over the set of all the joint actions. In their discussion it is necessary to assume that there are at least four players, and in some particular examples even three players suffice. Here we exhibit

a correlation mechanism for two players which uses a special channel of communication. In this channel the players can send two signals, say, A and B. In turn, the players receive a signal only if both players sent B; otherwise the channel remains silent.

For every correlation matrix with rational entries we exhibit a procedure that generates it and terminates with probability one. In case where the procedure can be repeated many times it turns out that this mechanism can be improved and become incentive compatible.

The paper ends with Section 8, in which the proofs are given. The effort is devoted mainly to showing that the correlation procedure is not manipulable.

2 The Model

a The One-Shot Game

The repeated game is an infinite repetition of a one-shot game, G , that consists of:

- (i) two finite sets of actions, Σ_1 and Σ_2 . Denote $\Sigma = \Sigma_1 \times \Sigma_2$;
- (ii) two payoff functions h_1, h_2 ; $h_i : \Sigma \rightarrow \mathbb{R}$. Denote $h = (h_1, h_2)$;
- (iii) two information functions ℓ_1, ℓ_2 defined on Σ , and ranged to L_1, L_2 , respectively.

The repeated game is a *symmetric S-T*³ *information game* if, for every $(a, b) \in \Sigma$, either

- (a) $\ell_1(a, b) = \ell_2(a, b) = (a, b)$ (in this case we will say that the information is *standard*); or
- (b) $\ell_1(a, b) = a$ and $\ell_2(a, b) = b$ (in this case we will say that the information is *trivial*).

In words, either both players are informed of the action combination just played or both are informed solely of their own actions. In view of the Dalkey theorem [D], the requirement that the players are informed of their own actions is superfluous. All the results below also hold if trivial information means a null signal.

The domain of the functions ℓ_1 and ℓ_2 can be extended to all the joint mixed actions. Let $\Delta(A)$ denote the set of all the probability distributions on A . ℓ_i can be extended in a natural way to $\Delta(\Sigma_1) \times \Delta(\Sigma_2)$ so as to attain values in $\Delta(L_i)$. Thus, $\ell_i(p, q)$ is the probability distribution on L_i induced by (p, q) .

³ for standard-trivial.

b The Repeated Game

(i) *Pure strategies.* A *pure strategy* of player i is a sequence $f = (f^1, f^2, \dots)$, where $f^t : L_i^{t-1} \rightarrow \Sigma_i$. L_i^{t-1} is the cartesian product of L_i with itself $t - 1$ times and it consists of all player i 's possible histories of length $t - 1$. If f and g are two pure strategies of player 1 and 2, respectively, then $x_i^t(f, g)$ will denote the payoff of player i at stage t when f and g are the strategies played.

(ii) A *mixed strategy* is a probability distribution over the set of all the pure strategies of the repeated game. Let σ_i be a mixed strategy of player i . Denote by $E_{\sigma_1, \sigma_2}(x_i^t)$ the expected payoff of player i at stage t where the expectation is taken with respect to the measure induced by (σ_1, σ_2) .

c Upper Nash Equilibria of the Repeated Game

Let σ_i be a mixed strategy of player i . $H_i^*(\sigma_1, \sigma_2)$ is defined as the limit of the means of player i 's expected payoff. Precisely, $H_i^*(\sigma_1, \sigma_2) = \lim(1/T) \sum_{t=1}^T E_{\sigma_1, \sigma_2}(x_i^t)$ if the limit exists. We will say that (σ_1, σ_2) is an *upper Nash equilibrium* if $H_i^*(\sigma_1, \sigma_2)$, $i = 1, 2$, is defined and if for any other pair of mixed strategies, say, $(\bar{\sigma}_1, \bar{\sigma}_2)$,

$$(i) \quad H_1^*(\sigma_1, \sigma_2) \geq \limsup_T (1/T) \sum_{t=1}^T E_{\bar{\sigma}_1, \sigma_2}(x_1^t), \text{ and}$$

$$(ii) \quad H_2^*(\sigma_1, \sigma_2) \geq \limsup_T (1/T) \sum_{t=1}^T E_{\sigma_1, \bar{\sigma}_2}(x_2^t).$$

Denote by UEP the set of all the payoffs $H^*(\sigma_1, \sigma_2) = (H_1^*(\sigma_1, \sigma_2), H_2^*(\sigma_1, \sigma_2))$, where (σ_1, σ_2) is an upper Nash equilibrium.

Two more Nash equilibria concepts will be defined here.

d Uniform Equilibrium

$\sigma = (\sigma_1, \sigma_2)$ is a *uniform equilibrium* if $H^*(\sigma)$ is well defined, and if for any $\epsilon > 0$ there is T such that for every $t \geq T$ σ induces an ϵ -Nash equilibrium in the t -fold repeated game (see [S]).

Denote the set of all uniform equilibrium payoffs by UNIF.

e Banach Limit

Let L be a Banach limit. σ is an L -equilibrium if for every strategy of player 1, say, $\bar{\sigma}_1$, the following holds: $L(\{H_1^n(\bar{\sigma}_1, \sigma_2)\}_n) \leq L(\{H_1^n(\sigma_1, \sigma_2)\}_n)$, and a similar requirement for player 2.

Denote all the L -equilibrium payoffs by BEP_L .

f Correlated Equilibrium in the Repeated Game

In order to introduce the correlated equilibrium we have to add another element to the game: a mediator. Before starting the game a mediator picks, according to a known distribution a pair of signals (α, β) . He informs player 1 of the outcome α and player 2 of β . Relying on the private information they received from the mediator, the players choose a pure strategy to be played in the game. This correlation procedure is termed as a *correlated equilibrium* if neither player can increase his payoff (upper limit of the expected averages) by choosing differently (his pure strategies), as a function of his private information.

Formally, a *correlated equilibrium* is a tuple $(\Omega_1 \times \Omega_2, \mu, f, g)$, where $(\Omega_1 \times \Omega_2, \mu)$ is a product probability space. f (resp., g) is a measurable function from Ω_1 (resp., Ω_2) to the set of player 1's (resp., player 2's) pure strategies, which satisfies

$$\lim (1/T) \sum_{t=1}^T E_{\mu, f, g}(x_1^t) \geq \limsup_T (1/T) \sum_{t=1}^T E_{\mu, \bar{f}, g}(x_1^t),$$

(resp. the same inequality replacing x_1^t with x_2^t , \bar{f} with f and g of the right side with \bar{g}) for any measurable function \bar{f} (resp. \bar{g}).

Denote by UCEP the set of all pairs of payoffs associated with correlated equilibria. Similarly to the definition of L -equilibrium, we can define L -correlated equilibrium and denote by CEP_L the set of all payoffs associated with such equilibria.

g Saturated Games

If the Nash equilibrium payoffs set of a game Γ coincides with the set of the correlated equilibrium payoffs of it, we say that Γ is *saturated*.

Some examples and remarks.

- (i) Any zero sum game is saturated.
- (ii) The Nash equilibrium payoffs set of a saturated game is convex.
- (iii) Generically, any finite and saturated game has a unique Nash equilibrium payoff.

- (iv) Any game has an extension which is saturated. Namely, it is possible to add a random signal which takes place before the game starts, in such a way that the new game is saturated (see [F]).
- (v) Any two player repeated game with observable payoffs is saturated. This is a corollary of two characterizations. The first is a characterization of the Nash equilibrium payoffs in repeated games with observable payoffs (see [L1]) and the second is a characterization of the correlated equilibrium payoffs in general repeated games with imperfect monitoring (see [L2]).

3 Example

Before proceeding to the general results of the paper, we exemplify the way the internal correlation works and a way to immunize against manipulations. The necessity of the latter is demonstrated by using the payoff matrix attached. All the relevant details about a repeated game with S - T information can be compressed into a bimatrix endowed with asterisks. An asterisk will stand for standard information. The following example is inspired by an example of Aumann [A1].

	L	M	R
T	0,0*	0,0	0,0
E	0,0	6,6	2,7
B	0,0	7,2	0,0

The information that the players receive is standard only when the actions played are T and L , otherwise the information is trivial. In other words, if player 1 played T and player 2 played L , both players are informed that (T,L) was played. On the other hand, if another action combination was played, then the players are informed solely of their own actions.

We will show that (5,5) can be sustained by an equilibrium. Notice that (5,5) is the payoff corresponding to the correlated equilibrium (of the one-shot game) which attaches probability 1/3 to each of the pairs (M,E), (M,B), (R,E).

If both players play (1/2, 1/2, 0) (i.e., player 1 and player 2 play T,E and L,M with probability one-half each, respectively), then the common distribution of their signals is given by:

	L*	L	M
T*	1/4	0	0
T	0	0	1/4
E	0	1/4	1/4

Notice that given that * was not observed, the common distribution is given by

	L	M
T	0	1/3
E	1/3	1/3

which is the normalized bottom-right sub-matrix of the above matrix. With probability 3/4 the players end up with a correlation matrix which can give them both, according to the following description, the payoff 5. In order to assure the payoff (5,5) player 1 should play E for a long time if at the first stage he played E and he should play B if at the first stage he played T and did not observe *. Player 2 should play M for a long period of time for M and R for L (without an *).

With probability 1/4 they will receive standard information after the first stage. It will be common knowledge, and in this case they can start the procedure from the beginning. Namely, they should play (1/2, 1/2, 0) and so on. However, this correlation procedure is not incentive compatible because both have an incentive to deviate. For instance, if player 1 instead of playing (1/2, 1/2, 0) plays (1,0,0), he affects the correlation matrix to his benefit. That is, if player 2 adheres to the above instructions, the result is that with probability 1/2 the players play (T,L) (and observe *) and with probability 1/2 they play (T,M) which leads later to the actions (B,M). By that deviation player 1 increases his payoff from 5 to 7 ($h(B,M) = (7,2)$).

The procedure just described can be improved and thereby made into an equilibrium. Notice that by the deviation of player 1 the probability of * has been increased from 1/4 to 1/2. Therefore, by replicating the same procedure many times and by performing a standard statistical test, player 2 can detect (with high probability) any major deviation of player 1. The more the procedure replicates, the higher the precision with which a player can detect an opponent's deviation. In this context of infinite horizon and no discount, the number of replications can be unlimited. Thus, the accuracy of the statistical tests employed can become arbitrarily high. Therefore, the payoff (5,5) can be sustained by an exact equilibrium and not merely by an ϵ -equilibrium.

4 The Main Theorem

a Some Notations

We will say that two actions $a, b \in \Sigma_1$ are *indistinguishable* (we denote it by $a \sim b$) if for any $c \in \Sigma_2$, $\ell_2(a, c) = \ell_2(b, c)$. In words, a and b are indistinguishable if player 2, no matter what he is playing, cannot distinguish between them. A similar definition holds for the actions of player 2.

Notice that this definition does not use the particular structure of ℓ_1, ℓ_2 specified in Section 2. In games with $S-T$ information, for instance, two actions are indistinguishable if by playing each one of them the player cannot get a standard information no matter what the other player is doing.

Recall that we denote by $\Delta(\Sigma)$ the set of all the probability distributions Q over Σ , i.e., for any $(a, b) \in \Sigma$, $Q_{a,b} \geq 0$ and $\sum_{(a,b) \in \Sigma} Q_{a,b} = 1$. We will define a subset \mathcal{B} of $\Delta(\Sigma)$ as follows:

$$\mathcal{B} = \{Q \in \Delta(\Sigma) \mid$$

- (i) for every two indistinguishable actions $a, a' \in \Sigma_1$

$$\sum_{b \in \Sigma_2} h_1(a, b)Q_{a,b} \geq \sum_{b \in \Sigma_2} h_1(a', b)Q_{a,b};$$

and

- (ii) for every two indistinguishable actions $b, b' \in \Sigma_2$

$$\sum_{a \in \Sigma_1} h_2(a, b)Q_{a,b} \geq \sum_{a \in \Sigma_1} h_2(a, b')Q_{a,b}.$$

In words, a probability distribution Q is in \mathcal{B} if every action $a \in \Sigma_1$ is a best response (versus the mixed action induced by Q given a) among all the actions a' that are indistinguishable from it, and a similar requirement for every $b \in \Sigma_2$. Thus, any profitable “deviation” from Q is detectable.

* We will say that the information structure is *completely trivial* if there is a player whose actions all are pairwise indistinguishable.

b The Main Theorem

In two player repeated games with a symmetric non-completely trivial S – T information.

$$\text{UEP} = \text{UNIF} = \text{BEP}_L = h(\mathcal{B}) \cap \text{IR} \text{ for any Banach limit } L,$$

where IR is the set of all the individually rational payoffs.

c Remarks

(i) The case where the information is completely trivial is a private case of semi-standard information, which was treated in [L1]. In that case, the equilibrium payoff set is the convex hull of all the Nash equilibrium payoffs of the one-shot game.

(ii) The set \mathcal{B} is convex and h is multilinear. Therefore $h(\mathcal{B})$ is also convex.

5 Previous Results

a Characterization of Correlated Equilibrium Payoffs

The set of all correlated equilibrium payoffs in general two player repeated games with imperfect monitoring was characterized in [L2]. A complete description of this result requires representation of notions which lie beyond the scope of this paper. However, in the case of S – T information, the indistinguishability notion suffices.

In games with S - T the characterization of the correlated equilibrium payoff set is (see [L2]):

$$\text{UCEP} = \text{CEP}_L = h(\mathcal{B}) \cap \text{IR}, \text{ for all Banach limits } L. \tag{5.1}$$

This characterization suggests that the correlated equilibrium payoffs of the repeated game are only those individually rational payoffs associated with correlated actions (Q) that are immunized against non-detectable and profitable deviations (i.e., $Q \in \mathcal{B}$). Combining (5.1) with the main theorem we obtain:

Corollary: In two player repeated games with a non-completely trivial and symmetric S - T information

$$\text{UEP} = \text{UCEP}.$$

In other words, symmetric S - T repeated games are saturated.

b Equilibrium Payoffs as Limits of Payoffs in G_n

Denote by G_n the n -fold repeated game of G . G_n is described by two sets of actions, Σ_1^n and Σ_2^n , two payoff functions, h_1^n and h_2^n , which are the *average* of all the stage payoffs, and two information functions, ℓ_1^n and ℓ_2^n .

We can treat G_n as we treated G , and define the relation \sim on Σ_i^n . Extending ℓ_i^n to $\Delta(\Sigma_i^n)$ we can extend \sim to mixed strategies as follows. Two mixed strategies, σ and σ' in $\Delta(\Sigma_i^n)$, are *indistinguishable* ($\sigma \sim \sigma'$) if for any $\tau \in \Delta(\Sigma_{3-i}^n)$, $\ell_{3-i}^n(\sigma, \tau) = \ell_{3-i}^n(\sigma', \tau)$. In words, σ and σ' are indistinguishable if they yield the same distribution on L_{3-i}^n for any mixed strategy τ . Define for any integer n and $\epsilon \geq 0$ the following sets

$$D_\epsilon^n(1) = \{(\sigma, \tau) \in \Delta(\Sigma_1^n) \times \Delta(\Sigma_2^n) \mid h_1^n(\sigma, \tau) \geq h_1^n(\sigma', \tau) - \epsilon \text{ for all } \sigma' \sim \sigma\} \tag{5.2a}$$

$$D_\epsilon^n(2) = \{(\sigma, \tau) \in \Delta(\Sigma_1^n) \times \Delta(\Sigma_2^n) \mid h_2^n(\sigma, \tau) \geq h_2^n(\sigma, \tau') - \epsilon \text{ for all } \tau' \sim \tau\} \tag{5.2b}$$

$$D_\epsilon^n = D_\epsilon^n(1) \cap D_\epsilon^n(2). \tag{5.2c}$$

In words, D_ϵ^n is the set of all the pairs (σ, τ) in which a player cannot gain by more than ϵ if he plays a strategy indistinguishable from his prescribed one.

Denote $M = \bigcap_{\epsilon > 0} cl \cup_n h^n(D_\epsilon^n)$, where cl is the closure operator and $h^n = (h_1^n, h_2^n)$. The result in [L3] asserts that

$$M \cap IR \subseteq F, \text{ for } F = \text{UEP, UNIF, BEP}_L. \tag{5.3}$$

Moreover, M is convex.

Obviously, $\text{UNIF, UEP} \subseteq \text{UCEP}$ and $\text{BEP}_L \subseteq \text{CEP}_L$. Thus, in order to prove the main theorem it is sufficient, by (5.1) and (5.3), to show that $h(\mathcal{B}) \cap IR$ is included in the left side of (5.3). This will be done by showing that any payoff in $h(\mathcal{B})$ can be approximated by payoffs in $h^n(D_\epsilon^n)$ for some n and arbitrarily small ϵ .

6 The Case of Discrete Equivalence Classes

Without restricting generality, $0 \leq h_i \leq 1, i = 1, 2$. For the main construction of the strategy we will need a sub-matrix of the pattern (---|---^*) . The existence of such a sub-matrix is ensured in cases of non-completely trivial information where each player has at least two indistinguishable actions. If, however, at least one of the players has no two such actions, then the proof of the theorem is easy, as shown by the following proposition.

Proposition 1: If at least one player has no two indistinguishable actions, then

$$\text{UEP} = \text{UNIF} = \text{BEP}_L = h(\mathcal{B}) \cap IR.$$

Proof: Without loss of generality, player 1 is the one having no two indistinguishable actions. We claim first that

$$\text{Conv } h(D_0^1(2)) \cap IR \subseteq F, F = \text{UEP, UNIF, BEP}_L. \tag{6.1}$$

Since player 1 has no two indistinguishable actions (see (5.2a)), $D_0^1(1) = \Delta(\Sigma_1) \times \Delta(\Sigma_2)$. Thus, $D_0^1 = D_0^1(1) \cap D_0^1(2) = D_0^1(2)$. Notice that $D_0^1(2) \subseteq D_\epsilon^n$ for every n and ϵ . Therefore, $h(D_0^1(2)) \subseteq M$. Since M is convex, $\text{Conv } h(D_0^1(2)) \subseteq M$. Hence, by (5.3) one obtains the desired inclusion.

Second, we prove that

$$h(\mathcal{B}) \subseteq \text{Conv } h(D_0^1(2)). \tag{6.2}$$

Let $Q \in \mathcal{B}$. Denote by Q_b the marginal distribution on Σ_2 . For every $b \in \Sigma$ satisfying $Q_b > 0$, Q induces a distribution, denoted $Q(\cdot | b)$, over Σ_1 in the following way. $Q(a|b) = Q_{a,b}/Q_b$. Notice that $h(Q) \in \text{Conv}\{h(Q(\cdot | b)) | b \in \Sigma_2 \text{ and } Q_b > 0\}$. Moreover, as $Q \in \mathcal{B}$, any pair $(Q(\cdot | b), b)$ satisfying $Q_b > 0$ is an element of $D_0^1(2)$. Hence, $h(Q) \in \text{Conv } h(D_0^1(2))$. Therefore, (6.2) is obtained.

The facts that $\text{UEP}, \text{UNIF} \subseteq \text{UCEP}, \text{BEP}_L \subseteq \text{CEP}_L$, (5.1), (6.1) and (6.2) complete the proof of the proposition. //

It is left to discuss the case where both players have at least two indistinguishable actions. In other words, there are at least two rows and two columns without any asterisks. We will demonstrate the main idea of the proof by an example.

7 An Illustration of the Proof by an Example

In this section all the main ideas and the terminology used in the next section are introduced. Consider the distribution Q :

	L	M	R
T	0	0	0
E	0	1/5	1/5
B	0	1/5	2/5

Fig. 1.

over a 3×3 matrix. We assume that $Q \in \mathcal{B}$ and that we are not in the situation of Proposition 1. Since the information is not completely trivial we can find a submatrix with the following pattern (* indicates standard information):

	A	B
A		
B		*

Fig. 2.

The objective is to define for a given $\epsilon > 0$ a joint strategy (σ^n, τ^n) in G_n , for n to be specified later. (σ^n, τ^n) will possess the following two properties: (i) $h^n(\sigma^n, \tau^n)$ is close to $h(Q)$ and (ii) $(\sigma^n, \tau^n) \in D_\epsilon^n$.

(σ^n, τ^n) will consist of three phases: the correlation phase, the master phase, and the report phase.

a The Correlation Phase

Define the 4×4 matrix ϕ as follows:

$$\phi = \begin{array}{|cc|cc|} \hline 1 & 0 & 1 & 0 \\ \hline 0 & 1 & 0 & 1 \\ \hline 1 & 0 & 1 & 1 \\ \hline 0 & 1 & 1 & 1 \\ \hline \end{array}$$

Fig. 3.

Let $x = y = (1/4, 1/4, 1/4, 1/4)$.

In ϕ there are 10 ones. Now define the matrix ψ by replacing each 1 in ϕ by $1/10$. We will consider ψ as a correlation matrix as follows. An entry in ψ is picked with the attached probability (i.e., either with probability $1/10$ or with probability 0). If an entry in one of the left sub-matrices of ψ was chosen, player 2 plays M and otherwise R . Similarly, if an entry in one of the top sub-matrices was chosen, player 1 plays E and otherwise B . Notice that ψ and Q induce the same distribution over the original matrix. However, ψ has another nice property. ψ satisfies

$$\psi_{ij} = c\phi_{ij}x_iy_j, \text{ and } \phi_{ij} \in \{0,1\}, \text{ where } c > 0,$$

which will be useful in the procedure that follows.

The *jointly controlled correlation* (JCC) procedure is the following: before starting the game each player randomly picks a number from $\{1, \dots, 4\}$ with probability $1/4$ each (this corresponds to x, y). Suppose that player 1 picked k and player 2 picked m . In the coming stages the players communicate, using the sub-matrix of Figure 2, as will be explained later. The goal of the communication is to disclose enough information (about the pair (k, m)), so that in case $\phi_{k,m}$ is 0, both players will know it. However, the procedure the players follow in order to disclose information

should be designed carefully. Should $\phi_{k,m} = 1$, a player would not be able to infer much data, from the information he obtains, about the identity of the choice made by his opponent. For instance, if player 1 picked 3 and player 2 picked 1 ($k = 3$ and $m = 1$), player 1 should be able to deduce that m is not 2 (because $\phi_{3,2} = 0$ and player 1 knows that he should have been informed of that, had m been equal to 2). But player 1 should not be able to infer anything beyond the fact that the choice of player 2 belongs to $\{1,3,4\}$. Knowing the probability according to which player 2 picked a number, and after applying a Bayesian updating, player 1 ascribes a probability of $1/3$ to each one of the possibilities $m = 1,3,4$.

How Players Communicate: Basically, players communicate by answering “Yes-No” questions. As was noted before, the answers a player gets should not be too informative. The way by which it is accomplished is the following. A player always answers to the “Yes-No” question asked, but his opponent does not necessarily hear the answer. Both players answer simultaneously. Both hear the answer only if both answer “Yes”, and both hear nothing if *one* of them answers “No”. Since a sub-matrix of the form of the one in Figure 2 exists, this pattern of communication is plausible. Notice that if a player plays *A* for “No” and *B* for “Yes”, the outcome is an asterisk only if both answers were “Yes”. Obviously, if a player played *B* and did not get an asterisk he deduces that his opponent’s answer was “No”. However, if a player played *A* he can infer nothing about the previous move of the other player.

Now we are ready to describe the procedure in detail. For any 0 entry of ϕ attach a stage (there are 6 such entries). Enumerate these entries by $1, \dots, 6$. At stage 1, using the sub-matrix of Figure 2, the players check whether the entry 1, say, (i_1, j_1) , was chosen. Player 1 plays *B* if he picked i_1 and *A* otherwise, and player 2 plays *B* if he picked j_1 and *A* otherwise. In other words, player 1 answers the question “Did you pick i_1 ?” and player 2 answers the question “Did you pick j_1 ?” If the answer is affirmative a player should play *B* and, otherwise, *A*. If the signal a player got in the first stage is * (standard), then he knows that $(k, m) = (i_1, j_1)$. In particular, $\phi_{k,m} = 0$, which is considered a failure of the procedure. Thus, it should start all over from the beginning. However, if the signal of the player at stage 1 is trivial, they proceed to stage 2 in which they eliminate the possibility of $(k, m) = (i_2, j_2)$. That is, player 1 answers (by playing *B* for “Yes” and *A* for “No”) the question “Did you pick i_2 ?” and player 2 answers the question “Did you pick j_2 ?” If the signal is * then it is a failure and the JCC should start over (this will be called later an *unsuccessful round* of the JCC). Otherwise, the third stage is devoted to the elimination of $(k, m) = (i_3, j_3)$, and so forth.

Suppose that the six first stages passed without any asterisk having been observed. This means that $\phi_{k,m} \neq 0$, which is considered a success of the procedure, and the JCC is completed.

In a case where a failure occurs, the players should ignore the previous outcomes and independently pick once again a number from $\{1,2,3,4\}$ at random. Then they should proceed by answering “Yes-No” questions regarding their choices and continue that way until the first success. In later referrals we call the last part of the JCC (in which six consecutive stages have been passed without any asterisk) the *successful round* of the JCC.

Notice that if both players adhere to the JCC, the procedure terminates with probability 1. Moreover, the successful round of the JCC induces the same distribution as ψ . For instance, if player 1 picked 3, he knows only that player 2 picked with probability $1/3$ any of the columns 1, 3 or 4.

The correlation phase by itself is manipulable. A player can deviate from the prescribed procedure in two ways. He can choose one of $\{1,2,3,4\}$ with a different distribution than $(1/4,1/4,1/4,1/4)$, or he can report on one choice while he actually chose another. The report phase is aimed at preventing such deviations.

Let t_ϵ be the earliest time before which the JCC terminates with probability of at least $1 - \epsilon/3$.

b The Master Phase

At the moment the correlation phase is finished, the master phase starts. Here, each player plays the action determined by the correlation phase for a long period of time—namely, for at least $3t_\epsilon/\epsilon$ stages. Player 1 plays E for an outcome (i.e., if player 1 picked at the successful round of the JCC a number corresponding to a column) in one of the top sub-matrices of ϕ , and plays B otherwise; player 2 plays M for an outcome in one of the left sub-matrices, and R otherwise.

Remark: After finishing the JCC, the players play according to the correlation matrix. Assume, for instance, that player 1 plays a . After playing one time according to the correlation matrix, the posterior probabilities that player 1 has about player 2's actions might be changed. As a result, some deviations of player 1 may become profitable. However, the posterior probabilities might be altered only if by playing a , there is a positive probability (according to the correlation matrix) for at least one $*$. Therefore, there is no other action that is indistinguishable from a , which means that player 1 cannot deviate to another action without changing the probabilities of player 2's signals. Thus, there is no new undetectable and profitable deviation.

Notice that we are relying here on the symmetry assumption: if one player gets $*$ the other one also gets it. Otherwise, it might be that the action a is a best response (among all . . .) a priori but not a posteriori. In particular, this means that there is a positive probability for player 1 to get an asterisk while playing a , and that this signal is invisible to player 2.

All of these attest to the fact that the particular structure of symmetric S - T information is needed, not merely for the existence of a submatrix of Figure 2, which enables the correlation phase to take place: this specific information structure is used also in the design of the master plan. In the latter, the players play over and over again according to the same correlation and, due to the symmetric S - T information, without impairing its effectiveness (i.e., the correlation remains incentive compatible).

c The Report Phase⁴

At this phase the players check each other to see if somebody has deviated in the correlation phase. Player 1 (resp. 2) has to report the row (resp. column) he picked in the successful round of the JCC (recall that we define a joint strategy in G_n , where n is still to be specified).

Reports of that kind were introduced in [L2] and [L3]. The way to report on a row or a column is to encode it by a string of A 's and B 's, and to play in the submatrix of Figure 2 accordingly. That is, at the first part of the report phase, player 1 plays B , and player 2 plays either A or B depending on the string encoding the column he wants to report on (as the one picked by him at the successful round of the JCC). In the second part of the report phase the roles of the players are exchanged. Player 2 plays B and player 1 plays either A or B so as to report his chosen row. That completes the description of the joint strategy denoted by (σ^n, τ^n) .

We would like to show roughly why (σ^n, τ^n) possesses the desired properties. Notice that the length of the report phase is $2\log_2 4 = 4$. (2 for two players, 4 for four rows and columns in ϕ .) Set $n = t_\epsilon + \lceil 3t_\epsilon/\epsilon \rceil + 4$. If both players do not deviate then the JCC terminates with probability of at least $1 - \epsilon/3$ before t_ϵ . Thus, with probability of at least $1 - \epsilon/3$ the expected payoff at any stage of the master phase is $h(Q)$. Recalling that $0 \leq h_i \leq 1$, we obtain $\|h^n(\sigma^n, \tau^n) - h(Q)\| \leq \epsilon/3 + (t_\epsilon + 4)/n \leq \epsilon$. To convince the reader that $(\sigma^n, \tau^n) \in D_\epsilon^n$ we should show that there is no way to deviate in the correlation phase without affecting the distribution of the signals.

To see it, let us return to the example. Player 1, for instance, can cheat in the correlation phase. He can pick the third row and report as if he picked the second one. For example, at the stage the players are supposed to check whether the entry (2,1) (namely, the second row and the first column) was chosen, player 1 plays B and reports as if he chose 2, even though he chose the third row. Similarly, player 1 can pick row 2 – in the stage the players are supposed to check if the entry (2,1) was picked – play A , as if he did not choose 2. By this particular method of cheating, player 1 shifts weight from one entry to another and causes the induced distribution to be

$$\psi' = \begin{array}{|c|cc|cc|} \hline & 1/10 & 0 & 1/10 & 0 \\ \hline & 1/10 & 1/10 & 0 & 1/10 \\ \hline & \hline & 0 & 0 & 1/10 & 1/10 \\ \hline & 0 & 1/10 & 1/10 & 1/10 \\ \hline \end{array}$$

⁴ I owe the idea presented here to S. Sorin.

rather than the original one. However, this kind of cheating can be detected at the report phase with a positive probability.

At the report phase, player 1, after picking the second row in the correlation phase, should choose to report either on row 2 or row 3. If he reports row 2 there is a positive probability that player 2 picked 1 and the corresponding entry was not eliminated (because the JCC terminated without any *). That is, in the corresponding stage to (2,1) he, player 2, did not get the signal that he was supposed to get, an *. At that moment, player 2 knows that player 1 cheated. If, on the other hand, player 1 reports 3, then player 2 can detect the cheating in case (which has positive probability) he chose the second column. This is so because, if the choice of player 2 was the second column and the choice of player 1 was the second row, then the JCC terminates successfully (by ψ' (2,2) is not eliminated). However, in the report phase, player 1 reports that he chose the third column (at that moment he does not know whether player 2 picked 1, 2, or 4). Since $\phi_{3,2} = 0$ and the pair (3,2) should have been eliminated during the correlation phase, player 2 deduces that player 1 cheated.

8 The Proof of the Theorem

Since (i) $\text{UNIF} \subseteq \text{UEP} \subseteq \text{UCEP}$, (ii) $\text{BEP}_L \subseteq \text{CEP}_L$, and in view of (5.1), it is sufficient to show that $h(\mathcal{B}) \cap \text{IR} \subseteq \text{UNIF}, \text{BEP}_L$.

Let $Q \in \mathcal{B}$. We will define here a joint strategy, (σ^n, τ^n) in G_n , consisting of three phases, where n is to be specified later. (σ^n, τ^n) is constructed in such a way that it yields a payoff (in G_n) close to the payoff corresponding to Q (this is proved in Proposition 2). Moreover, (σ^n, τ^n) is immunized (up to the order of ϵ) against undetectable deviations (see Proposition 3). Both propositions imply that $h(Q)$ is included in the left side of (5.3), which concludes the proof.

By Proposition 1 we can assume, without loss of generality, that each of the players has at least two indistinguishable actions. Thus, there exists a sub-matrix which has the pattern of the 2×2 matrix of Figure 2. We will refer to that sub-matrix as the communication matrix.

Let $\epsilon > 0$.

The Correlation Phase

Constructing the auxiliary matrix

We will take a matrix Q' with rational entries that satisfies

$$\|Q - Q'\| < \epsilon/3|\Sigma| \tag{8.1}$$

We can assume that all the entries have a common denominator. Thus, $Q' = (q_{ij}/c)$, where $q_{ij}, c \in \mathbb{N}$, and $\sum q_{ij} = c$. Let p be the l.c.d. of q_{ij} , and Q_{ij} be a $q_{ij} \times q_{ij}$ matrix of 1's. Define now for every pair (i, j) a matrix, P_{ij} , with 0-1 entries as follows. If $q_{ij} = 0$, set $P_{ij} = 0$. Otherwise, define

$$P_{ij} = \begin{array}{|c|c|} \hline Q_{ij} & 0 \\ \hline & \\ \hline & Q_{ij} \\ \hline & \\ \hline 0 & \\ \hline & Q_{ij} \\ \hline \end{array}$$

where the sub-matrix Q_{ij} appears p/q_{ij} times in P_{ij} . Now define the auxiliary matrix ϕ by replacing the entry q_{ij}/c in Q' by the matrix P_{ij} . That is, $\phi = (P_{ij})$. Notice that ϕ has $p|\Sigma_1|$ rows and $p|\Sigma_2|$ columns.

Denote by ψ the normalized ϕ (all the entries sum up to 1).

Example: Suppose that Q' is the lower right sub-matrix of Figure 1. p is equal to 2 and ϕ is the matrix of Figure 3.

The *jointly controlled correlation (JCC)* is a procedure in which a player picks with the same probability a row of ϕ (player 1) or a column of ϕ (player 2), and eliminates the occurrence of ϕ -zero entry in the same manner that was described in Section 6.

Precisely, let $N(\phi)$ be the set of all the zero entries of ϕ . I.e., $(u, v) \in N(\phi)$ iff $\phi_{u,v} = 0$. Set $n(\phi) = |N(\phi)|$, the number of the zero entries in ϕ . Let β be a one-to-one map from $N(\phi)$ to $\{1, \dots, n(\phi)\}$. That is, β is an enumeration of all the zero entries of ϕ . Player 1 (resp. 2) should play B of the communication matrix at the stage $\beta(u, v)$ if he picked u (resp. v) and should play A otherwise. If a player, during one of these stages, receives * the procedure starts over. That is, a player should pick, once again, a row or column independently of the previous outcomes and should report on it. With probability 1 there will be a successful JCC in which no * was observed.

In the sequel we will refer to $\beta(u, v)$ as the stage in which (u, v) is checked in order to be eliminated. Notice that if $\phi(u, v) = 1$, (u, v) is not eliminated.

Denote by t_ϵ the earliest time before which the JCC terminates with a probability of at least $1 - \epsilon/3$. Let $n = t_\epsilon + \lceil 3t_\epsilon/\epsilon \rceil + \lceil \log p|\Sigma_1| \rceil + \lceil \log p|\Sigma_2| \rceil + 2$.

The Master Phase

The players should play in this phase according to the outcome of the previous phase. If player 1 picked in the successful JCC the row number $p(a_1 - 1) + b_1$ with $1 \leq a_1 \leq |\Sigma_1|$ and $b_1 \leq p$, he should play constantly the action enumerated a_1 . Similarly, if player 2 picked the column number $p(a_2 - 1) + b_2$ with $1 \leq a_2 \leq |\Sigma_2|$ and $b_2 \leq p$, he should play constantly the action enumerated a_2 .

This phase lasts $n - (\text{the length of the correlation phase}) - \lceil \log p |\Sigma_1| \rceil - \lceil \log p |\Sigma_2| \rceil - 2$ stages. Notice that the length of the correlation phase is a random variable.

The Report Phase

It is possible to encode a number of rows and columns by strings, consisting of the letters A and B , of length $\lceil \log p |\Sigma_1| \rceil + 1$ and of length $\lceil \log p |\Sigma_2| \rceil + 1$, respectively. Player 1, first, plays A or B of the communication matrix according to the string encoding of the row picked in the successful JCC of the first phase. At the same time, player 2 plays B so as to receive the report from 1.

Afterwards, player 2 reports his chosen column. He plays A or B of the communication matrix according to the string encoding his chosen column, while player 1 plays B .

Proposition 2: There is a constant c_1 , independent of n , such that

$$\|h^n(\sigma^n, \tau^n) - h(Q)\| < c_1 \epsilon.$$

Proof: Observe that if both players adhere to the JCC, then each of the nonzero entries of ϕ is a plausible outcome of the procedure. Moreover, all these entries have the same probability of being the outcomes. Furthermore, if we identify (as is actually done in the master phase) the rows $p(a_1 - 1) + b_1$ with a_1 and the columns $p(a_2 - 1) + b_2$ with a_2 , ϕ induces the same probability distribution over $\Sigma_1 \times \Sigma_2$ as Q' does. Thus, after t_ϵ stages with probability of at least $1 - \epsilon/3$, the players play a correlated strategy distributed like Q' .

Hence (recall (8.1)),

$$\begin{aligned} \|h^n(\sigma^n, \tau^n) - h(Q)\| &\leq \|h^n(\sigma^n, \tau^n) - h(Q')\| + \\ \|h(Q) - h(Q')\| &\leq (1 - \epsilon/3)(t_\epsilon + \lceil \log p |\Sigma_1| \rceil + \lceil \log p |\Sigma_2| \rceil + 2)/n + \epsilon/3 + \\ \|h(Q) - h(Q')\| &\leq (1 - \epsilon/3)(\epsilon/3)c_1 + \epsilon/3 + \epsilon/3 |\Sigma| \leq c_1 \epsilon, \end{aligned}$$

for a certain constant c_1 . //

Proposition 3: $(\sigma^n, \tau^n) \in D_{c_2, \epsilon}^n$, for a certain constant c_2 , independent of n .

Proof: First step. We first show that any deviation in the correlation phase is detectable, i.e., any change in the way a player picks a row or a column of ϕ , as well as any change in the way a player reports on these, will change his opponent's signals distribution. We will show it for player 1 and a similar proof will work for player 2.

A strategy $\bar{\sigma}^n$ induces a distribution on the pure strategies. In particular, it induces a distribution on plays in the first round of the JCC. Assume that $\bar{\sigma}^n$ is indistinguishable from σ^n . We will show that $\bar{\sigma}^n$ induces with τ^n the distribution given by Q' , by constructing a matrix γ , a function α , and a probability distribution q over the rows of γ .

A *play* consists of: (i) the instruction of what action to take at any stage of the first round of the JCC, and (ii) the report to be delivered in the report phase. However, we assumed that $\bar{\sigma}^n \sim \sigma^n$ and thus any action to be taken (in the first round of the JCC) according to $\bar{\sigma}^n$ is either indistinguishable from A or from B . Thus, a play actually indicates whether to play A (or something equivalent, to which we refer as A) or B in each stage. In other words, it specifies whether to report "I picked u " or "I did not pick u " in the stage $\beta(u, v)$, where the possible occurrence of (u, v) is checked. We denote the set of plays by K and a generic play by k .

Each of those plays will stand for one row of γ . The matrix γ will have $p|\Sigma_2|$ columns: one for each column of ϕ . $\gamma(k, v)$ will be defined as u if according to the play k , player 1 reports "I picked u " in $\beta(u, v)$. It should be noted that if according to k player 1 reports successively "I picked u " and "I picked u' " in the respective stages $\beta(u, v)$ and $\beta(u', v)$, the choice of $\gamma(k, v)$ should be the first to be reported. The second one will not be observed. $\gamma(k, v) = 0$ if in the play k player 1 reports at the stage $\beta(u, v)$, "I did not pick u " for all u or if (u, v) is not checked at all.

Player 1 should report, in the report phase, which row from ϕ he picked. We denote the corresponding report of k by $\alpha(k)$. I.e., $\alpha(k)$ is a row of ϕ . Finally, q is the distribution on the plays induced by $\bar{\sigma}^n$. Thus, we have compressed the instruction of the strategy $\bar{\sigma}^n$ into γ, α and q . The instructions for how to play in the first round of the JCC are summarized by γ , and those for how to report in the report phase are summarized by α . q is the probability distribution according to which a specific play is chosen (if the strategy $\bar{\sigma}^n$ is followed). Notice that two different plays may appear similar in the matrix γ . In these cases they differ by their respective reports specified by α .

The forthcoming proof makes use of the following lemma which states that if $\bar{\sigma}^n$ and σ^n are indistinguishable, then both induce (together with τ^n) the same correlation matrix.

Lemma 1: Let γ be a matrix with $|K|$ rows and $p|\Sigma_2|$ columns, where $\gamma(k, v)$ is an integer between 0 and $p|\Sigma_1|$ for all k, v . Suppose, furthermore, that there is a function α from K to $\{1, \dots, p|\Sigma_1|\}$ identifying a row in γ with one in ϕ , and a probability distribution, say, q , assigning a positive probability to any row of γ . Assume:

- (1) if $\phi(u, v) = 0$ and $\alpha(k) = u$, then $\gamma(k, v) \neq 0$,
- (2) for every u, v , $q\{k; \gamma(k, v) = u\}$ is equal to $1/p|\Sigma_1|$ if $\phi(u, v) = 0$ and 0 otherwise,
- (3) $q\{k; \alpha(k) = u_0, \gamma(k, v) = 0\} / q\{k; \gamma(k, v) = 0\} = 1/\#\{u; \phi(u, v) = 1\}$ for all v and u_0 , satisfying $\phi(u_0, v) = 1$.

Then

- (a) $q\{k; \alpha(k) = u\} = 1/p|\Sigma_1|$ for every u .
- (b) If $\gamma(k, v) \neq 0$ and $\alpha(k) = u$, then $\phi(u, v) = 0$.

(b) can be rewritten as

$$(b') \alpha(k) = u, \phi(u, v) = 1 \implies \gamma(k, v) = 0.$$

The proof of the lemma is included in the Appendix.

In order to apply the lemma, we have to confirm that γ, q and α defined above satisfy (1)–(3) of Lemma 1. The assumption that $\bar{\sigma}^n \sim \sigma^n$ ensures that (1)–(3) are satisfied. Hypothesis (1) is satisfied because $\phi(u, v) = 0$ means that player 1 should report “I picked u ” at $\beta(u, v)$. There is a positive probability of column v to be picked by player 2. $\alpha(k) = u$ means that player 1 reports “I picked u ” at the report phase. However, this is inconsistent with $\gamma(k, v) = 0$ in the following sense: in a case where player 1 plays according to k and player 2 had picked v , the JCC will terminate successfully (because $\gamma(k, v) = 0$ means that while playing according to k player 1 reports at the stage $\beta(u, v)$, “I did not pick u ” for every u). $\alpha(k) = u$ means that player 1 reports, in the report phase, “I picked u ”. Knowing his chosen column, v , player 2 knows that if, indeed, player 1 chose u , he, player 2, should have heard the report, “I picked u ” in the appropriate stage (because $\phi(u, v) = 0$ means that the entry (u, v) should be eliminated). However, the actual report he heard was “I did not pick u ”, which is inconsistent. Since inconsistent messages are assigned zero probability we conclude that if, indeed, $\phi(u, v) = 0$ and $\alpha(k) = u$ then $\gamma(k, v)$ cannot be zero. Hence (1) is satisfied.

Hypothesis (2) is satisfied because $q\{k; \gamma(k, v) = u\}$ is the probability that, in the checking stage of the entry (u, v) an * will be observed, given that player 2 picked v . It should be equal to the probability assigned by the original (σ^n, τ^n) . That is, $1/p|\Sigma_1|$ if $\phi(u, v) = 0$ (i.e., the entry (u, v) is checked in order to be eliminated) and 0 if $\phi(u, v) = 1$ (i.e., (u, v) is not checked at all).

Hypothesis (3) is satisfied because, given that the JCC passed all the checkings without any *, the probability of getting a report “I picked u ” in the report phase (the left side of (3)) should equal the probability assigned by the original strategies (σ^n, τ^n) to the same event (i.e., the right side of (3)).

We can now apply Lemma 1. Let us convince the reader that conclusions (a) and (b) of the lemma mean that $(\bar{\sigma}^n, \tau^n)$ induces the same correlation matrix as (σ^n, τ^n) . I.e., both induce ψ . For this purpose we identify with u all the plays, k , for which $\alpha(k) = u$. Recall that there are $p|\Sigma_1|$ rows in ϕ .

Part (a) states that the total probability of those plays, k , identified with u , is $1/p|\Sigma_1|$, which is the probability of u (according to ψ). Now fix k satisfying $\alpha(k) = u$. (b') states that if $\phi(u, v) = 1$ (i.e., (u, v) is not supposed to be eliminated), then, indeed, $\gamma(k, v) = 0$. Namely, player 1, playing according to k , reports "I did not pick \bar{u} " in all the stages $\beta(\bar{u}, v)$. Thus, (u, v) is not eliminated. Furthermore, by hypothesis (1), if an entry (u, v) should be eliminated, (i.e., $\phi(u, v) = 0$), then it is eliminated also with $\bar{\sigma}^n$, because $\gamma(k, v) \neq 0$. From (b) and (1) we get that, given k satisfying $\alpha(k) = u$, the distribution over the v 's is the same as the distribution over the v 's induced by ψ , given the row u . This allows us to identify every play k that satisfies $\alpha(k) = u$ with u and to conclude that $(\bar{\sigma}^n, \tau^n)$ induces the distribution ψ .

To recapitulate, $(\bar{\sigma}^n, \tau^n)$ and (σ^n, τ^n) induce the same correlation if the first trial of the JCC is successful (i.e., if the successful round of the JCC is the first one). The same argument applies also to successful rounds which are not the first one (i.e., those that follow one or a few unsuccessful rounds). Thus, $\bar{\sigma}^n$, together with τ^n , induces the same correlation matrix as σ^n induces. Hence, we are in the context of the second step, where we assume that the correlation phase ends up with the correlation matrix Q' .

Second step: In view of the first step we may assume that the players follow the prescribed strategies in the correlation phase and, therefore, $(\bar{\sigma}^n, \tau^n)$ generate the distribution Q' for use in the master phase. It remains to be proven that given this assumption no player can gain much by a nondetectable (indistinguishable) deviation. We will prove that player 1 cannot gain in the master plan by more than $\epsilon/3$ by deviating to $\bar{\sigma}^n$ which agrees with σ^n on the correlation phase.

The notion of indistinguishable deviation in games with $S-T$ information means that at any stage and after any history the deviating strategy should assign to any equivalence class (induced by \sim) the same probability assigned by the prescribed strategy. Recall that $Q \in \mathcal{B}$ (i.e., any action assigned a positive probability by Q is a best response, within the equivalence class, to the expected action of the opponent). Since Q' is close to Q up to $\epsilon/3|\Sigma|$ any action assigned a positive probability by Q' is an $\epsilon/3$ best response, within the equivalence class, to the expected action (according to Q') of the opponent. For the sake of simplicity we divide the analysis into two cases.

Case 1: During the master phase player 1 does not get any additional information. This means that player 1 cannot update his beliefs about the action of player 2. Thus, player 1 expects player 2 to play the (mixed) action he (player 1) expected immediately after the correlation phase. However, his chosen action was then an $\epsilon/3$ best response (within the equivalence class) and it remains so during the master phase.

To recapitulate, in a case where no additional information is acquired during the master phase player 1 plays his $\epsilon/3$ best response (within the equivalence class). Since $\bar{\sigma}^n \sim \sigma^n$, by playing $\bar{\sigma}^n$ player 1 cannot gain by more than $\epsilon/3$ in the master phase.

Case II: During the master phase player 1 gets additional information. The master phase consists of repeating many times the action corresponding to the outcome of the correlation phase. In games with *S-T* information additional information can be acquired only by playing an action which is indistinguishable from any other action. I.e., an action, say, $a \in \Sigma_1$, for which there is an action $b \in \Sigma_2$ satisfying $\ell_1(a, b) = (a, b)$. We call such an action an *informative action*. However, any deviation *to* informative action (in order to obtain additional information) or *from* informative action (in the case when the outcome of the correlation phase corresponds to such an action) is indistinguishable. Thus, any indistinguishable deviation may yield at most a profit of $\epsilon/3$.

In the second step we have shown that $\bar{\sigma}^n$ does not increase the payoff of player 1 in the *master plan* by more than $\epsilon/3$. However, it may increase his payoff in the other phases. But the total length of these phases is negligible compared to the length of the master plan. Thus, we conclude that

$$\begin{aligned}
 h_1^n(\bar{\sigma}^n, \tau^n) &\leq h_1^n(\sigma^n, \tau^n) + \frac{(1 - \epsilon/3)t_\epsilon + \log p|\Sigma_1| + \log p|\Sigma_2| + 2}{n} + \epsilon/3 \\
 &\leq h_1(\sigma^n, \tau^n) + c_2 \epsilon, \text{ for some constant } c_2,
 \end{aligned}$$

which concludes the proof of Proposition 3. //

Propositions 2 and 3 imply that $h(Q) \in \bigcap_{\epsilon > 0} c! \cup_{n=1}^\infty h^n(D_\epsilon^n)$, which is, in view of (5.3), what is needed to prove the main theorem.

We conclude the paper by parenthetically noting that the correlation phase can be improved so as to hold not only for rational entries matrices but also for more general ones. However, for proving the main theorem, generating a rational correlation matrix (in the proof, Q') close enough to the matrix in question (Q) was sufficient.

Appendix

Proof of Lemma 1: Denote for any row u of ϕ , $\alpha^{-1}(u) = \{k \in K; \alpha(k) = u\}$. We will prove (a) first.

Fix u_0 . There are v 's for which $\phi(u_0, v) = 1$. Fix such v . By (3)

$$\begin{aligned}
 q\{\alpha^{-1}(u_0)\} &\geq q\{k; \alpha(k) = u_0, \gamma(k, v) = 0\} = \frac{q\{k; \gamma(k, v) = 0\}}{\#\{u; \phi(u, v) = 1\}} \\
 &= \frac{1 - q\{k; \gamma(k, v) \neq 0\}}{\#\{u; \phi(u, v) = 1\}}
 \end{aligned}$$

by (2)

$$= \frac{1 - \sum_{u:\phi(u,v)=0} 1/p|\Sigma_1| - \sum_{u:\phi(u,v)=1} q\{k;\gamma(k,v) = u\}}{\#\{u;\phi(u,v) = 1\}}$$

and again by (2)

$$= [1 - \frac{p|\Sigma_1| - \#\{u;\phi(u,v) = 1\}}{p|\Sigma_1|}] / \#\{u;\phi(u,v) = 1\} = 1/p|\Sigma_1|$$

Thus, $q\{\alpha^{-1}(u_0)\} \geq 1/p|\Sigma_1|$ for every u_0 . Since there are $p|\Sigma_1|$ such u_0 , we obtain (a).

We now proceed to the proof of (b). Fix v , a column in ϕ . Suppose that u satisfies $\phi(u, v) = 0$. By (1) we obtain $\gamma(k, v) \neq 0$ for every k satisfying $\alpha(k) = u$. Thus, by (2),

$$q\{k;\alpha(k) = u, \gamma(k, v) \neq 0\} = q\{k;\alpha(k) = u\} = 1/p|\Sigma_1|. \tag{A.1}$$

The last equality holds because of (a).

From (A.1) we get, by taking a union over all u such that $\phi(u, v) = 0$,

$$\begin{aligned} q\{k;\text{there is } u \text{ s.t. } \alpha(k) = u, \phi(u, v) = 0, \text{ and } \gamma(k, v) \neq 0\} \\ = \#\{u;\phi(u, v) = 0\}/p|\Sigma_1|. \end{aligned} \tag{A.2}$$

On the other hand, we obtain by (2) that

$$\begin{aligned} q\{k;\gamma(k, v) \neq 0\} &= \sum_u q\{k;\gamma(k, v) = u\} \\ &= \#\{u;\phi(u, v) = 0\}/p|\Sigma_1|. \end{aligned} \tag{A.3}$$

Thus, the probabilities in the left side of (A.2) and (A.3) coincide. Notice that the event of the left side of (A.2) is defined with the additional qualification that there is u s.t. $\alpha(k) = u$ and $\phi(u, v) = 0$, while that of (A.3) is defined without it. Since both probabilities: the one with the additional restriction and the one without it coincide, we conclude that $\gamma(k, v) \neq 0$ implies that there is u s.t. $\alpha(k) = u$ and $\phi(u, v) = 0$. Thus, if $\gamma(k, v) \neq 0$ and $\alpha(k) = u$, it must be that $\phi(u, v) = 0$, which concludes the proof. //

References

- [APS] Abreu D, Pearce D, Stachetti E (1986) Optimal Cartel Equilibria with Imperfect Monitoring, *Journal of Economic Theory*, 39, 251-269.
- [A1] Aumann RJ (1974) Subjectivity and Correlation in Randomized Strategies, *Journal of Mathematical Economics*, 1, 67-95.
- [A2] Aumann RJ (1981) Survey of Repeated Games, in: *Essays in Game Theory and Mathematical Economics in Honor of Oskar Morgenstern*. Bibliographisches Institut, Mannheim-Wien-Zürich, 11-42.
- [B] Barany I (1987) Fair Distribution Protocols on How the Players Replace Fortune, CORE Discussion Paper No. 8718, CORE, Universite Catholique de Louvain. To appear in MOR.
- [D] Dalkey N (1953) Equivalence of Information Patterns and Essentiality Determinant Games, in: Kuhn HW, Tucker AW (eds.), *Contributions to the Theory of Games II*, Princeton University Press, Princeton, N.J. and *Annals of Mathematical Studies*, 28, 217-243.
- [F] Forges F (1990) Universal Mechanisms, *Econometrica*, 58, 6, 1341-1364.
- [FL] Fudenberg D, Levine D (1989) An Approximate Folk Theorem with Imperfect Private Information, mimeo.
- [FM] Fudenberg D, Maskin D (1986) Discounted Games with Unobservable Actions, I: One-Sided Moral Hazard, mimeo.
- [K] Kohlberg E (1975) Optimal Strategies in Repeated Games with Incomplete Information, *International Journal of Game Theory*, 4, 1-24.
- [L1] Lehrer E (1986) Two Player Repeated Games with Non-Observable Actions and Observable Payoffs, to appear in *Mathematics of Operations Research*.
- [L2] Lehrer E (1987) Correlated Equilibria in Two-Player Repeated Games with Non-Observable Actions, to appear in *Mathematics of Operations Research*.
- [L3] Lehrer E (1989) On the Equilibrium Payoff Set in Two Player Repeated Games with Imperfect Monitoring, mimeo.
- [L4] Lehrer E (1989) Lower Equilibrium Payoffs in Repeated Games with Non-Observable Actions, *International Journal of Game Theory*, 18, 57-89.
- [L5] Lehrer E (1990) Nash Equilibria of n -Player Repeated Games with Semi-Standard Information, *International Journal of Game Theory*, 19, 191-217.
- [M] Mertens J-F, Sorin S, Zamir S (1990) Repeated Games, to be published.
- [R] Radner R (1986) Repeated Partnership Games with Imperfect Monitoring and No Discounting, *Review of Economic Studies*, 53, 1, 43-58.
- [RY] Rubinstein A, Yaari M (1983) Repeated Insurance Contracts and Moral Hazard, *Journal of Economic Theory*, 30, 74-97
- [S] Sorin S (1988) Supergames, to appear in *Game Theory and Applications (the First OSU Conference)*, Academic Press.

Received: October 1989

Revised version: September 1990

Final version: January 1991