Correlated Equilibria, Incomplete Information and Coalitional Deviations*

Francis Bloch, Bhaskar Dutta †

September 23, 2008

Abstract

This paper proposes new concepts of strong and coalition-proof correlated equilibria where agents form coalitions at the interim stage and share information about their recommendations in a credible way. When players deviate at the interim stage, coalition-proof correlated equilibria may fail to exist for two-player games. However, coalition-proof correlated equilibria always exist in dominance-solvable games and in games with positive externalities and binary actions. *JEL Classification Numbers: C72*

Keywords: correlated equilibrium, coalitions, information sharing, games with positive externalities

*We thank an Associate Editor and two anonymous referees for very helpful comments and suggestions on an earlier draft of the paper. Dutta also gratefully acknowledges support from ESRC Grant RES-000-22-0341.

†Bloch: corresponding author, Ecole Polytechnique, Route de Saclay, 91128 Palaiseau, France, tel:+33 1 69 33 30 45, francis.bloch@polytechnique.edu. Dutta: Department of Economics, University of Warwick, Coventry CV4 7AL, England, B.Dutta@warwick.ac.uk

1 Introduction

A game with *communication* arises when players have the opportunity to communicate with each other prior to the choice of actions in the actual game. The presence of a *mediator* is a particularly powerful device in such games because it allows players to use *correlated* strategies - the mediator (privately) recommends actions to each player according to the realization of an agreed upon correlation device.

A *correlated equilibrium* is a *self-enforcing* correlated strategy profile because no individual has an incentive to deviate from the recommendation received by her, given the information at her disposal. However, if players can communicate with each other, it is natural to ask whether coalitions of players cannot exchange information about the recommendations received by them and plan mutually beneficial *joint* deviations.

An important constraint on possible joint deviations is that the sharing of information must satisfy a "credibility" constraint. We borrow concepts of credible information-sharing from the literature on cooperative game theory with incomplete information[1] to develop two refinements of correlated equilibria. The first concept is analogous to that of strong Nash equilibrium. A correlated strategy profile is a *strong correlated equilibrium* if it is immune to deviations by coalitions of essentially myopic players who do not anticipate any further deviations after the coalition has implemented its blocking plan. The second concept is that of *coalition-proof correlated equilibrium.* According to this concept, coalitions take into account the possibility that sub-coalitions may enforce further deviations.

Notice that in our framework, coalitions plan deviations at the *interim* stage - that is, after the mediator has communicated his recommendation to each player. Einy and Peleg (1995) also define interim notions of strong and coalition-proof correlated equilibrium.[2] Of course, coalitions could also form at the *ex ante* stage, that is before the mediator has communicated his recommendations to the players. Moreno and Wooders (1996), Milgrom and Roberts (1996), Ray (1996) focus on these ex ante concepts. In section 3, we construct examples which illustrate the differences between our solution concept and these alternative definitions. In this section, we also show that if the action sets of all individuals are restricted to *two* identical actions, then the positive externality games studied by Konishi et al (1997) have interim correlated coalition-proof equilibria.

[1] The classic reference in this literature is Wilson (1978). The current paper adapts the notion of the credible core of Dutta and Vohra (2005) to this setting.

[2] See also Ray (1998).

2 Definitions

Let N denote the set of players, indexed by $i = 1, 2, .., n$. Each player has a finite set of pure strategies, A_i, with generic element a_i. Let $A = \prod_{i \in N} A_i$ denote the Cartesian product of the individual strategy sets. The utility of player i is given by $u_i : A \to \Re$

A *correlated strategy* μ is a probability distribution over A.

We consider an extended game with a mediator. Before the game is played, a mediator privately sends recommendations to the players, a, chosen according to the probability distribution μ. Each player observes his recommendation and then proceeds to playing the game.

Consider a correlated strategy μ and coalition S. Suppose members of S have privately received the recommendations a_S. How can they plan mutually beneficial deviations from μ? Any plans to "block" μ must depend on their beliefs about the realization of μ. Moreover, each individual i's belief about the realization of μ depends upon the recommendation received by i himself as well as the information about a_S which can be *credibly* transmitted by members of S to each other. In what follows, we adapt the notion of the *credible core* of Dutta and Vohra (2005) to this setting.

Suppose all members of S *believe* that the recommendations received by the players lie in some subset E of A. We will call such a set E an *admissible event*, and describe some restrictions which must be satisfied by such an event. First, an element a_{-S} can be ruled out only by using the private information of members of S. Since we will use conditional expected utilities to evaluate action profiles, we can without loss of generality express this requirement as $E = E_S \times A_{-S}$. Second, if $i \in S$, then her claim that she has not received recommendation a'_i cannot depend on the claims made by other members of S. Hence, E_S must be the cartesian product of some sets $\{E_i\}_{i \in S}$. Third, no agent can, after receiving her own recommendation, rule out the possibility that the "true" profile of recommendations lies in the set E. Hence, an admissible set for the coalition S, must satisfy the following.

Definition 1 *Given μ, an event E is admissible for S if and only if*

$$E = \prod_{i \in S} E_i \times A_{-S}, \text{ and } \sum_{\widetilde{a_{-i}} \in E_{-i}} \mu(a_i, \widetilde{a_{-i}}) > 0 \text{ for all } i \in S, a_i \in E_i,$$

where $E_{-i} = \prod_{j \in S \setminus i} E_j \times A_{-S}$.

Given an admissible event E, we define player i's conditional probability of a_{-i} given a_i and E as

$$\widetilde{\mu}(a_{-i}|a_i, E_{-i}) = \frac{\mu(a_{-i}, a_i)}{\sum\limits_{\widetilde{a_{-i}} \in E_{-i}} \mu(a_i, \widetilde{a_{-i}})}.$$

We also define the marginal probability over a_{-S} given a_i and E as:

$$\overline{\mu}(a_{-S}|a_i, E_{-i}) = \sum_{\alpha_{S\setminus i} \in \Pi_{j\in S\setminus i} E_j} \widetilde{\mu}(a_{-S}, \alpha_{S\setminus i}|a_i, E_{-i}).$$

A *blocking plan* for coalition S, η_S, is a correlated strategy over A_S. Thus, $\eta_S(a_S)$ denotes the probability with which any $a_S \in A_S$ is played once the blocking plan is implemented.

Once the blocking plan is implemented, a player i in S who has received the recommendation a_i^* has the following posterior belief over the actions in the game:

$$\gamma_i(a) \equiv \gamma_i(a_{-S}, a_S) = \overline{\mu}(a_{-S}|a_i^*, E_{-i})\eta_S(a_S).$$

Given a_i and E, player i evaluates the correlated strategy μ according to:

$$U_i(\mu|a_i, E_{-i}) = \sum_{a_{-i} \in E_{-i}} \widetilde{\mu}(a_{-i}|a_i, E_{-i})u_i(a_i, a_{-i}).$$

Player i evaluates the blocking plan according to:

$$U_i(\eta_S|a_i, E_{-i}) = \sum_{\widetilde{a} \in A} \gamma_i(\widetilde{a})u_i(\widetilde{a}).$$

Definition 1 ensures that if members of S each claim to have received recommendations in the set E_i, then no individual in S can conclude that some individual has lied given knowledge of his own recommendation. However, this condition by itself does not guarantee that each individual in S will believe the claims of other members of S. We explain below why there should be some further restriction on an admissible event before individuals can agree on a plan to block a correlated strategy μ.

Suppose E is an admissible event for coalition S, and $i \in S$. We want to rule out the possibility that i, after receiving the recommendation $a_i' \notin E_i$ actually claims to have received a recommendation in E_i. We first define the set of recommendations that player i *could have claimed to receive.*

Let $V_i(E) = \{a_i' \in A_i \setminus E_i |$ there is $a_{-i} \in E_{-i}$ such that $\mu(a_{-i}, a_i') > 0\}$.

If player if i receives $a_i' \in V_i(E)$, there is a positive probability for the event E_{-i}. Hence player i might gain by lying to the other coalition members and claiming that he received a recommendation in E_i.

For any coalition S, a blocking plan η_S on an admissible event E satisfies *self selection* w.r.t. the correlated strategy profile μ if for all $i \in S$ and all $a'_i \in V_i(E)$,

$$U_i(\mu|a'_i, E_{-i}) \geq U_i(\eta_S|a'_i, E_{-i}) \quad (1)$$

So, if this equation is satisfied, and if i has actually received the recommendation a'_i in $V_i(E)$, then her expected payoff from μ is at least as high as that from η_S. This implies that she has no incentive to agree to the blocking plan. Thus, her agreement to the blocking plan η_S signals that she has indeed received a recommendation in E_i.

Definition 2 *A coalition S blocks the correlated strategy μ if there exists a blocking plan η_S and admissible event E such that*

(i) η_S satisfies self-selection on E w.r.t. μ.

(ii) For all $i \in S$, for all $a_i \in E_i$, $U_i(\mu|a_i, E_{-i}) < U_i(\eta_S|a_i, E_{-i})$.

The underlying idea behind this notion of blocking is the following. If members of a coalition agree to a blocking plan, this information should be used to update players' information over the recommendations received by other players in the coalition. So, E defines the event for which all players in S have an incentive to accept the blocking plan η_S. Every player in S thus updates his beliefs by assuming that players in $S\backslash i$ have received recommendations in $\prod_{j \in S\backslash i} E_j$. If given these updated beliefs, all players in S have an incentive to accept the blocking plan η_S, then the coalition S blocks the correlated strategy over the event E.[3]

Definition 3 *A correlated strategy μ is an interim strong correlated equilibrium (ISCE) if there exists no coalition S that blocks μ.*

If the coalition S is a singleton, $E_{-i} = A_{-i}$ and the self-selection constraint is vacuous. A singleton coalition $\{i\}$ thus blocks the correlated strategy μ for the event $\{a_i\}$ if there exists a mixed strategy σ_i such that

[3]The Associate Editor has pointed out that Definition 2 restricts the scope for blocking compared to the corresponding definition in Dutta and Vohra (2005) because it requires the coalition S to use the same blocking plan η_S for every realization $a_S \in E_S$. However, we have deliberately adopted this restrictive notion of blocking. If a coalition is to use a possibly different blocking plan for each realization $a_S \in E_S$, then the members in S have to communicate the original recommendation received by them to a "new" mediator. This raises complicated issues of incentive compatibility, particularly when we use the same ideas to define coalition-proofness. Our more restrictive definition avoids these problems.

$U_i(\mu|a_i, A_{-i}) < U_i(\sigma_i|a_i, A_{-i})$. Hence, for singleton coalitions, our definition corresponds to the usual definition of correlated equilibrium.

As with the concept of strong Nash equilibrium, the concept of interim strong correlation equilibrium implicitly assumes that players are *myopic* when they plan deviations. Alternatively, they can sign *binding* agreements to enforce a blocking plan. That is why we do not impose any incentive compatibility constraints on deviations. In particular, the deviating coalition does not take into account the possibility that there may be further deviations.

Following Bernheim, Whinston and Peleg (1987), we now define a notion of coalition-proof equilibrium when coalitions form *after* players have received recommendations from the mediator. Notice that if a nested sequence of coalitions each form blocking plans, then the posterior beliefs of players "later on" in the sequence keep changing. Suppose, for instance, that the original correlated strategy is μ, and coalition S^1 considers a blocking plan η_1 on the admissible event E^1. Then, players in S^1 believe that the recommendations sent by the mediator lie in the set E^1. Moreover, the posterior beliefs of players in S^1 are different from their prior beliefs. Now, consider "stage 2" when the coalition $S^2 \subset S^1$ contemplates a blocking plan η_2 on the admissible event E^2. First, their prior beliefs coincide with the posterior beliefs formed at the end of stage 1. Second, players in S^2 now believe that the mediator has recommended an action profile in E^2. Implementation of the blocking plan η_2 will result in a new set of posterior beliefs for players in S^2, and this change in posterior beliefs will also change the way in which players evaluate blocking plans. This needs to be kept in mind when defining an interim notion of coalition-proofness, and also provides the motivation for the following definitions.

Consider a coalition $S \subseteq N$, and a *blocking sequence* $\mathbf{B} = \{(S^k, \eta_k, E^k)\}_{k=1}^K$ to the correlated strategy μ, where

(i) $S^1 \equiv S$, and for each $k = 2, \ldots, K$, $S^k \subset S^{k-1}$.

(ii) For each $k > 1$, $E_i^k \subset E_i^{k-1}$ for $i \in S^k$, and $E_i^k = E_i^{k-1}$ for $i \notin S^k$.

(iii) Each η_k is a correlated strategy over A_{S^k}.

A blocking sequence thus consists of a nested sequence of coalitions, blocking plans for every coalition in the sequence, and events at every step of the sequence, which satisfy the natural conditions that they become finer for players who belong to successively smaller coalitions. We now define the posterior beliefs for each coalition in the blocking sequence, in order to define admissible events and self-selection.

At the initial step, let $\gamma^0(a) \equiv \mu(a)$ for all $a \in A$. Choose any $k \in \{1, \ldots, K\}$, and $i \in S^k$. Then,

$$\widetilde{\mu}^k(a_{-i}|a_i, E^k_{-i}) = \frac{\gamma^{k-1}(a_{-i}, a_i)}{\sum\limits_{\widetilde{a_{-i}} \in E^k_{-i}} \gamma^{k-1}(a_i, \widetilde{a_{-i}})}.$$

Similarly, the marginal probability over a_{-S^k} given a_i and E is:

$$\overline{\mu}^k(a_{-S^k}|a_i, E^k_{-i}) = \sum_{\alpha_{S^k\backslash i} \in \Pi_{j \in S^k \backslash i} E^k_j} \widetilde{\mu}^k(a_{-S^k}, \alpha_{S^k\backslash i}|a_i, E^k_{-i}).$$

Once the blocking plan η^k is implemented, a player i in S^k who has received the recommendation a^*_i has the following posterior belief over the actions in the game:

$$\gamma^k_i(a) = \overline{\mu}^k(a_{-S^k}|a^*_i, E^k_{-i})\eta_{S^k}(a_{S^k}).$$

These definitions allow us to define recursively the beliefs γ^k of members of coalitions along the blocking sequence. We now use this sequence of beliefs to define credible deviations.

Consider any coalition $S \subseteq N$. A *blocking sequence* $\mathbf{B} = \{(S^k, \eta_k, E^k)\}^K_{k=1}$ to μ is *legitimate* if

(i)$\forall\ k = 1, \ldots, K,$ given γ^{k-1}, E^k is admissible for S^k and

(ii) η_k satisfies self-selection $w.r.t. \gamma^{k-1}$

In order to define the concept of interim coalition-proof correlated equilibrium (ICPCE), we first define the notion of self-enforcing blocking plans.

Definition 4 *Let S be any coalition.*

(i) If $|S| = 1$ (say $S = \{i\}$) , then any mixed strategy σ_i is a self-enforcing blocking plan against any correlated strategy.

(ii) Recursively, suppose that self-enforcing blocking plans have been defined for all coalitions of size smaller than $|S|$ against any correlated strategy. Then, S has a self-enforcing blocking plan η_1 against the correlated strategy μ if

(a) There is an admissible event E^1 for S given μ such that η_1 satisfies self-selection on E^1 w.r.t. μ, and

(b) There is no legitimate blocking sequence $\{(S, \eta_1, E^1), (S^2, \eta_2, E^2)\}$ where η_2 is a self-enforcing blocking plan for S^2 w.r.t. γ^1 such that

$$\sum_{a_{-i} \in E^2} \widetilde{\mu}^2(a_{-i}|a_i, E^2_{-i})u_i(a_i, a_{-i}) < \sum_{a \in A} \gamma^2_i(a)u_i(a) \text{ for all } i \in S^2, \text{ for all } a_i \in E^2_i.$$

Definition 4 is a direct transposition to our context of the definition of *self-enforcing deviations* in Bernheim, Whinston and Peleg (1987). In our setting, we need to keep track of events and beliefs in order to check whether an initial blocking plan is immune to further deviations by subcoalitions, making the definition considerably more notation intensive than the original definition. Notice also that, while we defined blocking sequences of arbitrary length, our definition of self-enforcing blocking plans only uses blocking sequences of length two. This is due to the fact that, once we recursively define the set of self-enforcing blocking plans for smaller coalitions, we only need to check that a blocking plan is immune to one-step deviations by subcoalitions using self-enforcing blocking plans.

We may now define interim coalition-proof correlated equilibrium (ICPCE) using the same steps as the definition of interim strong correlated equilibria.

Definition 5 *A coalition S blocks the correlated strategy μ with a self-enforcing blocking plan η_S if there exists an admissible event E such that*

(i) η_S satisfies self-selection on E w.r.t. μ.

(ii) For all $i \in S$, for all $a_i \in E_i$, $U_i(\mu|a_i, E_{-i}) < U_i(\eta_S|a_i, E_{-i})$.

Definition 6 *A correlated strategy μ is an interim coalition proof correlated equilibrium (ICPCE) if there exists no coalition S that blocks μ with a self-enforcing blocking plan.*

Some remarks are in order. First, any strong correlated equilibrium is a coalition-proof correlated equilibrium, as any self-enforcing blocking plan is a blocking plan. Second, because any blocking plan by a single player coalition is self-enforcing, any coalition-proof correlated equilibrium is a correlated equilibrium.

3 Discussion

Different definitions of strong and coalition proof correlated equilibria have already been proposed in the literature. Moreno and Wooders (1996) and Milgrom and Roberts (1996) consider coalitional deviations at the *ex ante* stage, before agents have received their recommendations.[4] Formally, in

[4]Ray (1996) also proposes a notion of coalition proof correlated equilibrium at the ex ante stage. Intuitively, his concept differs from Moreno and Wooders (1996)'s, Milgrom and Roberts (1996)'s and ours in that deviating coalitions cannot choose a new correlation device, but must abide by the fixed correlation device of the extended game.

their setting, a blocking plan is a mapping η_S from A_S to ΔA_S, assigning a correlated strategy over A_S to any possible recommendation a_S. Players evaluate the correlated strategy μ according to the expected utility

$$U_i(\mu) = \sum_{a \in A} \mu(a) u_i(a).$$

Given a blocking plan η_S against the correlated strategy μ, the induced distribution over actions is given by

$$\widehat{\mu}(a) = \sum_{\alpha_S \in A_S} \mu(\alpha_S, a_{-S}) \eta_S(a_S | \alpha_S)$$

and players evaluate the blocking plan according to

$$U_i(\eta_S) = \sum_{a \in A} \widehat{\mu}(a) u_i(a)$$

Definition 7 *A correlated strategy μ is an ex ante strong correlated equilibrium (ESCE) if there exists no coalition S and blocking plan η_S such that $U_i(\eta_S) > U_i(\mu)$ for all $i \in S$.*

As above, we define self-enforcing ex ante blocking plans recursively. Any blocking plan by a one-player coalition is self-enforcing. Given that self-enforcing blocking plans have been defined for all coalitions T with $|T| < |S|$, a blocking plan η_S generating a distribution $\widehat{\mu}$ is self-enforcing, if there exists no coalition $T \subset S$, and self-enforcing blocking plan η_T for T generating a distribution $\widehat{\mu_T}$ such that $U_i(\widehat{\mu}_T) > U_i(\widehat{\mu})$ for all i in T.

Definition 8 *A correlated strategy μ is a ex ante coalition proof correlated equilibrium (ECPCE) if there is no coalition S and self-enforcing blocking plan η_S such that $U_i(\eta_S) > U_i(\mu)$ for all $i \in S$.*

Coalitional incentives to block at the ex ante and interim stage cannot be compared. On the one hand, it may be easier for coalitions to block at the ex ante stage. Consider for example a correlated strategy in a two-player game putting equal weight on two outcomes with payoffs $(0, 3)$ and $(3, 0)$. At the ex ante stage, this correlated strategy has expected value 1.5 for every player, and would be blocked by another outcome with payoffs $(2, 2)$. However, at the interim stage, neither of the two realizations can be blocked by both players. On the other hand, coalitions may find it easier to block at the interim stage, when a correlated strategy puts weight on an outcome with very low payoffs for the players. The following example illustrates this point.[5] This example also highlights another important difference between ICPCE and ECPCE - the former may fail to exist in two-person games where ECPCE always exist.[6]

Example 1 *Consider a two-player game where player* 1 *chooses the row and player* 2 *the column.*

	b_1	b_2	b_3
a_1	$4, 4$	$-4, 0$	$0, 4.1$
a_2	$1, 1$	$1, 1$	$-1, 0$
a_3	$0, 0$	$0, -1$	$2, 2$

This game possesses two pure strategy Nash equilibria (a_2, b_2) and (a_3, b_3). Both equilibria are dominated by a correlated equilibrium putting weight 1/2 on (a_1, b_1), and 1/4 on (a_2, b_1) and (a_2, b_2). To check that this correlated strategy forms an equilibrium, observe that player 1's best response is to play a_1 when player 2 plays b_1 and a_2 when she believes that player 2 plays b_1 and b_2 with equal probability. Similarly, player 2's best response is to play b_1 when she believes that player 1 has received recommendation a_1 with probability 2/3 and recommendation a_2 with probability 1/3, and to play b_2 when player 1 plays a_2.

Next, we observe that any correlated equilibrium putting positive weight on the cell (a_2, b_2) is dominated by the pure strategy Nash equilibrium (a_3, b_3). When instructed to play a_2 and b_2, both players obtain a payoff less or equal to 1. Hence, irrespective of their beliefs on the recommendation received by the other player, both players have an incentive to switch to the profile (a_3, b_3).

Finally, we show that there is no other correlated equilibrium in the game. If the correlated strategy puts zero weight on cell (a_2, b_2), then strategy a_2 is strictly dominated for player 1 and strategies b_1 and b_2 are strictly dominated for player 2. As a correlated equilibrium cannot put weight on strictly dominated strategies, it must concentrate all the weight on strategy b_3. Hence, if a correlated equilibrium does not put weight on cell (a_2, b_2), it must concentrate all the probability on the cell (a_3, b_3).

[5] This example considers a two-player game with three strategies per player. We believe – but have not formally shown – that it is the minimum number of strategies needed to construct such an example

[6] Moreno and Wooders (1996) show that the set of ECPCE is equal to the set of ex ante Pareto-undominated correlated equilibria in two-person games.

Einy and Peleg (1995) define an interim notion of strong and coalition proof correlated equilibrium. Their concept differs from ours in two important respects. First, they assume that members of a blocking coalition freely share information about their recommendations.[7] Second, they assume that a coalition blocks if all its members are made better off *for any realization* of the initial correlated strategy. Formally, they define a blocking plan as a mapping from A_S (the set of recommended strategies in μ) to ΔA_S. In their equilibrium concept, a coalition S blocks, if for all possible realizations a_S, the blocking plan is a strict improvement for all players in S.

There is no inclusion relation between the set of strong (and coalition proof) correlated equilibria defined by Einy and Peleg (1995) and the set of strong (and coalition-proof) correlated equilibria defined in this paper. On the one hand, the fact that members can freely share information about their recommendations makes deviation easier in Einy and Peleg (1995)'s sense. On the other hand, their – very strong – requirement that coalitional members are better off for any realization of the correlated strategy makes deviations harder. Consider for instance the following example of a three-player game due to Einy and Peleg (1995).[8]

Example 2 *(Einy and Peleg (1995)) Consider the following three-player game, where player I chooses rows* (a_1, a_2)*, player II chooses columns* (b_1, b_2) *and player III chooses matrices* (c_1, c_2)*.*

	b_1	b_2
a_1	3,2,0	0,0,0
a_2	2,0,3	2,0,3

c_1

	b_1	b_2
a_1	3,2,0	0,3,2
a_2	0,0,0	0,3,2

c_2

Einy and Peleg (1995) argue that the following is a strong correlated equilibrium.

	b_1	b_2
a_1	1/3	0
a_2	0	1/3

c_1

	b_1	b_2
a_1	0	1/3
a_2	0	0

c_2

[7]In the context of exchange economies with private information, this is equivalent to the notion of "fine" core proposed by Wilson (1978). The problem of course is that players' announcements about the recommendation they received is not verifiable, and blocking plans may not be credible in our sense.

[8]This example was also used by Moreno and Wooders (1996) to show that a three-player game may fail to admit an ECPCE.

To prove their claim, they note that for any two-player coalition, there exists one realization of the correlated strategies for which no strict improvement is possible. (For $S = \{1, 2\}$, the realization (a_1, b_1), for $S = \{2, 3\}$, the realization (b_2, c_2) and for $S = \{1, 3\}$, the realization (a_2, c_1).)

With our definition, we claim that this correlated strategy is *not* a ISCE. Consider the coalition $S = \{1, 2\}$ and the realization (a_2, b_2). Player 1 then knows that 3 has received the recommendation c_1 and that 2 has received the recommendation b_2. So, player 1 expects a payoff of 2 if he follows the recommendation. Player 2 puts equal probability to (a_2, c_1) and (a_1, c_2), and so expects a payoff of 1.5.

Consider the admissible event $E_1 = \{a_1, a_2\}, E_2 = \{b_2\}$. For this event, both players have a blocking plan (a_1, b_1) which satisfies self-selection on E. To see this, notice that $V_1(E)$ is empty since $E_1 = A_1$. However, $V_2(E) = \{b_1\}$.

Hence, coalition $\{1, 2\}$ blocks the correlated strategy at the realization (a_2, b_2) and the correlated strategy is not a strong correlated equilibrium.

We now provide some remarks on the existence of ISCE and ICPCE. Clearly, ISCE is a very demanding concept, and one does not expect ISCE to exist for general classes of game. Since the definition requires deviations to be self-enforcing, the existence of ICPCE is easier to guarantee. We first observe that the main existence theorem of Moreno and Wooders (1996) can be adapted to our context to provide a sufficient condition for existence. Moreno and Wooders (1996) note that if there exists a correlated strategy which Pareto-dominates any other action profile in the set of actions surviving iterated elimination of dominated strategies, this correlated equilibrium is an ECPCE (Corollary p. 92). In our context, ex ante Pareto-dominance is not sufficient to guarantee existence, but if there exists a *pure strategy profile* which Pareto-dominates any other pure strategy profile in the set of strategies surviving iterated elimination of dominated strategies, it forms an ICPCE of the game.[9] Hence, any dominance-solvable game admits an ICPCE. Furthermore, as noted by Milgrom and Roberts (1996), in games with strategic complementarities admitting a unique Nash equilibrium, or for which utilities are monotonic in the actions of the other players, there exists a pure action profile which dominates any other action profile in the set of strategies surviving iterated elimination of dominated strategies (Theorem 2, p. 124), so that ICPCE exist.

[9]The proof of this statement, which is omitted, follows exactly the same steps as the proof in Moreno and Wooders (1996).

Finally, ICPCE are shown to exist in a class of games with positive externalities studied by Konishi et al. (1997), for which strong Nash equilibria always exist. Players all have the same action set $A = \{a_1, a_2\}$ and have utility functions which depend on the action they choose and the set of players who choose the same action, and are increasing in the set of players choosing the same action. This situation corresponds for example to the games of standardization in the presence of network externalities studied by Farrell and Saloner (1985) and Katz and Shapiro (1985). Formally, for any player i, $u_i(\mathbf{a}) = u_i(a^i, A^i)$, where a^i denotes the action chosen by player i, and A^i denotes the set of players choosing action a^i. Moreover, $u_i(a^i, A^i) \geq u_i(a^i, A'^i)$ if $A'^i \subseteq A^i$. That is, utility is nondecreasing if a larger set of individuals uses the same action that i uses.

Proposition 1 *Any game with binary actions and positive externalities admits an interim coalition proof correlated equilibrium.*

Proof. From Konishi et al. (1997), we know that the game admits pure strategy strong Nash equilibria. Pick one of these strong Nash equilibria characterized by a partitioning of the agents, $\{B_1, B_2\}$, where B_i denotes the set of players choosing the action a_i for $i = 1, 2$. Moreover, choose the partition so that if B'_1 is a superset of B_1, then $\{B'_1, N \setminus B'_1\}$ is not a strong Nash equilibrium.

Let T be a coalition which has a profitable blocking plan η_T against the pure strategy recommendation which results in the partition $\{B_1, B_2\}$. We first claim that $T \cap B_i \neq \emptyset$ for $i = 1, 2$ – the deviating coalition must involve players from both coalitions. Suppose by contradiction that $T \subset B_1$. (A similar argument would hold if $T \subset B_2$). Because $\{B_1, B_2\}$ is a strong Nash equilibrium, there must exist an agent $i \in T$ for whom $u_i(a_1, B_1) \geq u_i(a_2, B_2 \cup T)$. But if $T \subset B_1$, then for any outcome $\{C_1, C_2\}$ in the support of the blocking plan, $C_1 \subseteq B_1$. Hence, for all outcomes in the support of the blocking plan, agent i either chooses action a_1 in a group containing $C_1 \subseteq B_1$ agents, or chooses action a_2 in a group containing $C_2 \subseteq B_2 \cup T$ agents. In either case, his utility is less than or equal to $u_i(a_1, B_1)$ and he cannot participate in the blocking plan.

Let $T_1 = T \cap B_1$ and $T_2 = T \cap B_2$. Consider the partition $\{B_1 \cup T_2, B_2 \setminus T_2\}$. By assumption, this partition is not a strong Nash equilibrium, and so there exists a deviating coalition S. We show that $S \subset T_2$. First notice that $B_1 \cap S = \emptyset$. If members of B_1 had an incentive to deviate collectively in the partition $\{B_1 \cup T_2, B_2 \setminus T_2\}$, they would also have an incentive to deviate in the partition $\{B_1, B_2\}$, contradicting the fact that $\{B_1, B_2\}$ is a strong

Nash equilibrium. Notice furthermore that if there exists a deviating coalition S containing members of T_2 and $B_2 \setminus T_2$, then there also exists another deviating coalition S' only containing members of T_2. Hence, if there is no deviating coalition S satisfying $S \subset T_2$, it must be that all deviating coalitions are included in $B_2 \setminus T_2$. Consider then the largest deviating coalition, S, for which $u_i(a_1, B_1 \cup T_2 \cup S) > u_i(a_2, B_2 \setminus T_2)$ for all $i \in S$, and the resulting partition $\{B_1 \cup T_2 \cup S, B_2 \setminus (T_2 \cup S)\}$. Again, this partition is not a strong Nash equilibrium, and there must exist a deviating coalition U. By the same argument as above, we must have $U \subset B_2 \setminus (T_2 \cup S)$. The process can be repeated until the formation of the partition $\{N, \emptyset\}$, at which point we reach a contradiction, because this partition is not a strong Nash equilibrium, and it is impossible to construct a deviating coalition. Hence, there must exist a deviating coalition S from $\{B_1 \cup T_2, B_2 \setminus T_2\}$ such that $S \subset T_2$.

Finally, we show that this implies that there exists a self-enforcing blocking plan, η_S against the original deviation η_T. Consider the plan where members of S always choose action a_2. Every member i of S will then receive at least $u_i(a_2, (B_2 \cup S) \setminus T_2)$ after deviating. By sticking to the recommendation a_1, he would receive at most $u_i(a_1, (B_1 \cup T_2 \setminus S) \cup \{i\}) \leq u_i(a_1, B_1 \cup T_2)$. Because S is a deviating coalition from the partition $\{B_1 \cup T_2, B_2 \setminus T_2\}$, $u_i(a_2, B_2 \cup S \setminus T_2) > u_i(a_1, B_1 \cup T_2)$ for all $i \in S$, and hence the blocking plan η_S is profitable. Finally, the blocking plan is self-enforcing, because no subcoalition of S can guarantee a higher payoff to all its members, as this would involve some players moving back to action a_1. ■

To conclude, we view our study of coalitional deviations in games with communication as a first step towards the study of coalitional deviations in general Bayesian games. Our definition of credible information sharing could easily be adapted to a setting where agents have different (privately known) types, and our equilibrium concepts could easily be applied to general games with incomplete information. We plan to pursue this agenda in future research, thereby making progress on the study of cooperation and coalition formation among agents with incomplete information.

4 References

Bernheim, D. , B. Peleg and M. Whinston (1987) "Coalition-proof Nash Equilibria: Concepts", *J. Econ. Theory* **42**, 1-12.

Dutta, B. and R. Vohra (2005) "Incomplete Information, Credibility and the Core", *Math. Soc. Sci.* **50**, 148-165.

Einy, E. and B. Peleg (1995) "Coalition Proof Communication Equilibria",

in *Social Choice, Welfare and Ethics,* (W. Barnett, H. Moulin, M. Salles and N. Schofield, eds.), Cambridge, New York and Melbourne: Cambridge University Press.
Farrell, J. and G. Saloner (1985) "Standardization, Compatibility and Innovation", *RAND J. Econ.* **16**, 70-83.
Katz, M. and C. Shapiro (1986) "Technology Adoption in the Presence of Network Externalities", *J. Polit. Economy* **94**, 822-841.
Konishi, H. , M. Le Breton and S. Weber (1997) "Pure Strategy Nash Equilibria in Group Formation Games with Positive Externalities", *Games Econ. Behav.* **21**, 161-182.
Milgrom, P. and J. Roberts (1996) "Coalition Proofness and Correlation with Arbitrary Communication Possibilities", *Games Econ. Behav.* **17**, 113-128.
Moreno, D. and J. Wooders (1996) "Coalition Proof Equilibrium", *Games Econ. Behav.* **17**, 80-112.
Ray, I. (1996) "Coalition Proof Correlated Equilibrium: A Definition", *Games Econ. Behav.* **17**, 56-79.
Ray, I. (1998) "Correlated Equilibrium as a Stable Standard of Behavior", *Rev. Econ. Design*, **3**, 257-269.
Wilson, R. (1978) "Information, Efficiency and the Core of an Economy", *Econometrica* **46**, 807-816.