The Geometry of Nash Equilibria and Correlated Equilibria and a Generalization of Zero-Sum Games *

Yannick Viossat†

Stockholm School of Economics and
Laboratoire d'économétrie, Ecole polytechnique, Paris

Abstract

A pure strategy is coherent if it is played with positive probability in at least one correlated equilibrium. A game is pre-tight if in every correlated equilibrium, all incentives constraints for non deviating to a coherent strategy are tight. We show that there exists a Nash equilibrium in the relative interior of the correlated equilibrium polytope if and only if the game is pre-tight. Furthermore, the class of pre-tight games is shown to include and generalize the class of two-player zero-sum games.

Keywords: correlated equilibrium; zero-sum games; dual reduction

1 Introduction

The set of correlated equilibria of a finite game is a polytope containing the Nash equilibria. A better understanding of the location of the Nash

*This is a thoroughly revised version of "Geometry, correlated equilibria and zero-sum games", cahier du laboratoire d'économétrie 2003-032, Ecole polytechnique, Paris

†E-mail: yannick.viossat@polytechnique.org. This paper was completed while at the Stockholm School of Economics. I have since moved to Dauphine University, Paris, France. My current address is: CEREMADE, Université Paris-Dauphine, Place du Maréchal de Lattre de Tassigny, 75775 Paris cedex 16, France

equilibria within this polytope might allow not only to shed light on the connections between Nash equilibria and correlated equilibria, but also to design more efficient algorithms to compute Nash equilibria. The question was first studied by Evangelista and Raghavan (1996) [6] and by Gomez Canovas et al. (1999) [9]. They showed that in bimatrix games, extreme Nash equilibria are extreme points of the correlated equilibrium polytope. More recently, Nau et al. (2004) [17, proposition 2] proved the following result, which applies to finite games with any number of players. Call coherent the pure strategies that are played with positive probability in at least one correlated equilibrium. If there is a Nash equilibrium in the relative interior of the correlated equilibrium polytope, then:

(i) The Nash equilibrium assigns positive probability to every coherent strategy of every player;

(ii) In every correlated equilibrium, the incentive constraints for non deviating from one coherent strategy to another coherent strategy are all satisfied with equality.

In particular, if condition (ii) is not satisfied, then all Nash equilibria belong to the relative boundary of the correlated equilibrium polytope.

This leaves several questions unanswered: is it possible to find necessary *and sufficient* conditions for a Nash equilibrium to lie in the relative interior of the correlated equilibrium polytope? If so, is it possible to check that these conditions are satisfied without computing the correlated equilibria of the game? Finally, are these conditions satisfied by many games and in conceptually important classes of games?

This article answers these questions positively: first, condition (ii) is actually necessary and sufficient. Thus, there exists a Nash equilibrium in the relative interior of the correlated equilibrium polytope *if and only if* condition (ii) is satisfied. Second, it is possible to check that a game satisfies (ii) without computing its correlated equilibria. Third, the class of games satisfying (ii), which we call pre-tight games, has positive measure. Furthermore, in the two-player case, it includes and generalizes the class of two-player zero-sum games. In particular, Nash equilibria are exchangeable, Nash equilibrium payoffs and correlated equilibrium payoffs coincide, and profiles of correlated equilibria's marginals are Nash equilibria. Up to our knowledge, this is the largest class of games in which it is known that Nash equilibria are exchangeable.

Several proofs are based on dual reduction (Myerson, 1997 [16]). An additional interest of this paper is thus to illustrate the use of this technique to investigate the properties of correlated equilibria and Nash equilibria.

The material is organized as follows: the next section is devoted to basic notations and definitions. In section 3, we recall the definition of tight games (Nitzan, 2005 [19]) and introduce the class of pre-tight games. Section 4 shows that whether a game is tight (resp. pre-tight) or not may be checked without computing its correlated equilibria. The link between tight and pre-tight games is made precise in section 5. Topological properties of the sets of tight and pre-tight games are studied in section 6. In section 7, we show that the relative interior of the correlated equilibrium polytope contains a Nash equilibrium iff the game is pre-tight. Finally, in section 8, we show that in the two-player case, pre-tight games include and generalize zero-sum games.

2 Notations and definitions

Let

$$G = \{I, (S_i)_{i\in I}, (U_i)_{i\in I}\}$$

denote a finite game in strategic form; I is the nonempty finite set of players, S_i the nonempty finite set of pure strategies of player i and $U_i : \times_{i\in I} S_i \to \mathbb{R}$ the utility function of player i. Let $S := \times_{i\in I} S_i$ and $S_{-i} := \times_{j\in I\setminus\{i\}} S_j$. For any finite set Σ, $\Delta(\Sigma)$ denotes the set of probability distributions over Σ. As usual, letting $s \in S$ and $t_i \in S_i$, we denote by (t_i, s_{-i}) the strategy profile that differs from s only in that its $i-$component is t_i. Similarly, for any mixed strategy profile $\sigma \in \times_{i\in I}\Delta(S_i)$, we may write $\sigma = (\sigma_i, \sigma_{-i})$.

A *correlated strategy* of the players in I is a probability distribution over the set S of pure strategy profiles. Thus $\mu = (\mu(s))_{s\in S}$ is a correlated strategy if:

$$\text{(nonnegativity constraints)} \qquad \mu(s) \geq 0 \quad \forall s \in S \tag{2.1}$$

$$\text{(normalization constraint)} \qquad \sum_{s\in S} \mu(s) = 1 \tag{2.2}$$

For every $\mu \in \Delta(S)$, we let $U_i(\mu) := \sum_{s\in S} \mu(s)U_i(s)$. Furthermore, for every pure strategy s_i in S_i, $\mu(s_i \times S_{-i}) := \sum_{s_{-i}\in S_{-i}} \mu(s)$ denotes the marginal probability of s_i in μ. We say that the pure strategy $s_i \in S_i$ is *played* in μ if this marginal probability is positive. The correlated strategy of the other players given s_i is then denoted by $\mu(.|s_i) \in \Delta(S_{-i})$:

$$\forall s_{-i} \in S_{-i}, \mu(s_{-i}|s_i) = \frac{\mu(s)}{\mu(s_i \times S_{-i})}$$

For every (i, s_i, t_i) in $I \times S_i \times S_i$, let h_{s_i,t_i} denote the linear form on $\mathbb{R}^S$ which maps μ to

$$h_{s_i,t_i}(\mu) := \sum_{s_{-i}\in S_{-i}} \mu(s)[U_i(s) - U_i(t_i, s_{-i})]$$

Note for later purposes that:

Remark 2.1 *If $\mu(s_i \times S_{-i}) = 0$, then $h_{s_i,t_i}(\mu) = 0$ for all t_i in S_i.*

A correlated strategy μ is a *correlated equilibrium* (Aumann [2]) if

$$\text{(incentive constraints)} \quad h_{s_i,t_i}(\mu) \geq 0 \quad \forall i \in I, \forall s_i \in S_i, \forall t_i \in S_i \tag{2.3}$$

A possible interpretation is as follows: assume that before play a mediator chooses a strategy profile s with probability $\mu(s)$ and privately recommends s_i to player i. The incentive constraints (2.3) stipulate that if all the players other than i follow the recommendations of the mediator, then player i has no incentives to deviate from s_i to some other strategy t_i.

Since conditions (2.1), (2.2) and (2.3) are all linear in μ, the set of correlated equilibria is a polytope, which we denote by C. Furthermore, assimilating $\times_{i\in I} S_i$ to a subset of $\Delta(S)$, it is easily checked that the Nash equilibria are exactly the correlated equilibria μ with a product distribution; that is, such that

$$\forall s \in S, \mu(s) = \prod_{i\in I} \mu(s_i \times S_{-i})$$

The set of Nash equilibria is thus the intersection of the correlated equilibrium polytope and of the variety of product distributions.

We now introduce the classes of games that will be studied throughout.

3 Tight and pre-tight games

3.1 Tight games

Definition 3.1 *A game is* tight *(Nitzan [19]) if in any correlated equilibrium all the incentive constraints are tight. Formally,*

$$\forall \mu \in C, \forall i \in I, \forall s_i \in S_i, \forall t_i \in S_i, h_{s_i,t_i}(\mu) = 0 \tag{3.1}$$

This means that whenever a pure strategy s_i is played in a correlated equilibrium μ, every pure strategy of player i is a best-response to $\mu(\cdot|s_i)$.

Example 3.2

$$G_1 = \begin{pmatrix} 1,-1 & -1,1 \\ -1,1 & 1,-1 \end{pmatrix} \qquad G_2 = \begin{pmatrix} 1,-1 & -1,1 & 0,-1 \\ -1,1 & 1,-1 & 0,-1 \end{pmatrix}$$

The game G_1 (i.e. Matching Pennies) is tight. Indeed, it has a unique correlated equilibrium: the Nash equilibrium σ in which both players play $(1/2, 1/2)$. Since σ is a completely mixed Nash equilibrium, it follows that in σ, all incentive constraints are satisfied with equality, hence (3.1) is satisfied.

By contrast, the game G_2 is not tight. Indeed, the mixed strategy profile in which the row player plays $(\frac{1}{2}, \frac{1}{2})$ and the column player $(\frac{1}{2}, \frac{1}{2}, 0)$ is a Nash equilibrium, hence a correlated equilibrium. However, the incentive constraint stipulating that player 2 has no incentive to deviate from his first strategy to his third is satisfied with strict inequality. Another way to see that G_2 is not tight is to note that the third strategy of player 2 is strictly dominated. Indeed:

Proposition 3.3 *If there exists a pure or mixed strategy which is strictly dominated then the game is not tight.*

Proof. Let σ be a Nash equilibrium. If the game is tight, then every pure strategy, hence also every mixed strategy of player i is a best-response to σ_{-i}. It follows that no mixed strategy is strictly dominated. ■

More examples will be given in section 4.5

3.2 Pre-tight games

Definition 3.4 *(Nau et al. [17]) The pure strategy s_i (resp. the pure strategy profile s) is* coherent *if it is played in correlated equilibrium; that is, if there exists a correlated equilibrium μ such that $\mu(s_i \times S_{-i}) > 0$ (resp. $\mu(s) > 0$).*

Denote by S_i^c the set of coherent pure strategies of player i.

Definition 3.5 *A game is* pre-tight *if in any correlated equilibrium all the incentive constraints for non deviating to a* coherent *strategy are tight. Formally,*

$$\forall \mu \in C, \forall i \in I, \forall s_i \in S_i, \forall t_i \in S_i^c, h_{s_i,t_i}(\mu) = 0 \tag{3.2}$$

Note that, by remark 2.1, (3.2) is equivalent to

$$\forall \mu \in C, \forall i \in I, \forall s_i \in S_i^c, \forall t_i \in S_i^c, h_{s_i,t_i}(\mu) = 0 \tag{3.3}$$

A game is thus pre-tight if, whenever a pure strategy s_i is played in a correlated equilibrium μ, every *coherent* pure strategy of player i is a best-response to $\mu(\cdot|s_i)$. This does not imply that every coherent pure strategy is played in all correlated equilibria. For instance, the game G_3 (below, left) is pre-tight, as follows from proposition 3.8. Furthermore, since the correlated strategy μ (below, center) is a completely mixed Nash equilibrium, it follows that every pure strategy is coherent. Nevertheless, the third column is not played in the Nash equilibrium ν (below, right).

$$G_3 = \begin{pmatrix} 1,-1 & -1,1 & 0,0 \\ -1,1 & 1,-1 & 0,0 \end{pmatrix} \quad \mu = \begin{array}{|c|c|c|} \hline 1/6 & 1/6 & 1/6 \\ \hline 1/6 & 1/6 & 1/6 \\ \hline \end{array} \quad \nu = \begin{array}{|c|c|c|} \hline 1/4 & 1/4 & 0 \\ \hline 1/4 & 1/4 & 0 \\ \hline \end{array}$$

Proposition 3.6 *Any tight game is pre-tight.*

Proof. Condition (3.2) is weaker than (3.1). ■

Proposition 3.7 *Any game with a unique correlated equilibrium is pre-tight.*

Proof. If a game has a unique correlated equilibrium σ, then σ is a Nash equilibrium. Furthermore, the set of coherent strategies of player i is simply the support of σ_i and for every pure strategy s_i in the support of σ_i, $\sigma(\cdot|s_i) = \sigma_{-i}$. Therefore, the game is pre-tight iff for every i in I, every pure strategy in the support of σ_i is a best-response to σ_{-i}. Since σ is a Nash equilibrium, this condition is satisfied. ■

Proposition 3.8 *Any two-player zero-sum game is pre-tight.*

This will be proved in section 8, proposition 8.1.

3.3 Best-Response Equivalence

Consider two games G and G' with the same sets of players and strategies, but with different utility functions:

Definition 3.9 *The games G and G' are* best-response equivalent *if for every player i in I, every pure strategy s_i in S_i, and every correlated strategy ν in $\Delta(S_{-i})$, the pure strategy s_i is a best-response to ν in G iff s_i is a best-response to μ_{-i} in G'.*

Proposition 3.10 *If G is tight (resp. pre-tight) then any game that is best-response equivalent to G is tight (resp. pre-tight).*

Proof. In the definitions of tight and pre-tight games, the utility functions only intervene via best-responses to correlated strategies of the other players. The result follows. ■

In particular, whether a game is tight or not (resp. pre-tight or not) is unaffected by positive affine transformations of the payoff functions. We now provide a criterion that allows to check that a game is tight (resp. pre-tight) without having to compute its correlated equilibria.

4 Characterization of tight and pre-tight games

The results are stated in section 4.1 and proved in sections 4.2 and 4.3. It is also shown that every pre-tight game has a quasi-strict Nash equilibrium (section 4.4), paving the way for section 7. Finally, examples of applications of the characterization of tight and pre-tight games are given in section 4.5.

4.1 Statement of the results

We first need a definition: for each player i in I, let α_i be a transition probability over the set of pure strategies of player i:

$$\begin{array}{rrcl} \alpha_i : & S_i & \rightarrow & \Delta(S_i) \\ & s_i & \rightarrow & \alpha_i * s_i \end{array}$$

Let $\alpha = (\alpha_i)_{i \in I}$ and for every strategy profile s in S, let

$$f(s,\alpha) := \sum_{i \in I} [U_i(\alpha_i * s_i, s_{-i}) - U_i(s)] \tag{4.1}$$

Definition *The vector of transition probabilities* α *is a* dual vector *(Myerson [16]) if for every* s *in* S*,* $f(s,\alpha) \geq 0$

Note that there always exists a dual vector: just take $\alpha_i * s_i = s_i$ for all i and s_i.

Proposition 4.1 *A game is tight iff there exists a dual vector* α *such that, for every player* i *in* I *and every pure strategy* s_i *in* S_i*, the mixed strategy* $\alpha_i * s_i$ *is completely mixed.*

Proposition 4.2 *A game is pre-tight iff there exists a dual vector* α*, and, for every player* i *in* I*, a subset* $S'_i \subseteq S_i$ *of pure strategies such that:*

(A) *For every player* i *in* I *and every pure strategy* s_i *in* S'_i*, the mixed strategy* $\alpha_i * s_i$ *has support* S'_i*.*

(B) *For every pure strategy profile* s *in* S *that does not belong to* $S' := \times_{i \in I} S'_i$*, we have* $f(s,\alpha) > 0$

In that case, S'_i *is the set of coherent pure strategies of player* i*. That is,* $S'_i = S^c_i$*.*

4.2 Proof of proposition 4.1

The proof relies on the strong complementary property of dual linear programs. We first need a definition and a few lemmas.

Definition (Myerson [16]) *Let s_i and t_i be two pure strategies of player i. The strategy t_i* jeopardizes *s_i if*

$$\forall \mu \in C, h_{s_i,t_i}(\mu) = 0$$

That is, t_i jeopardizes s_i if for every correlated equilibrium μ in which s_i is played, the pure strategy t_i is a best response to $\mu(.|s_i)$. Note that the definitions of tight and pre-tight games may be rephrased in terms of jeopardization:

Lemma 4.3 *A game is tight iff for every i in I, any pure strategy of player i jeopardizes all his pure strategies.*

Indeed the above condition is exactly:

$$\forall i \in I, \forall t_i \in S_i, \forall s_i \in S_i, \forall \mu \in C, h_{s_i,t_i}(\mu) = 0$$

which is equivalent to (3.1). Similarly, it follows from (3.3) that:

Lemma 4.4 *A game is pre-tight iff for every i in I, every coherent pure strategy of player i jeopardizes all his coherent pure strategies.*

Moreover, dual vectors arise as the solutions of a dual linear program, such that the solutions of the primal are the correlated equilibria of the game (Myerson [16]). Exploiting this fact and the strong complementary property of dual linear programs allows to show that:

Lemma 4.5 (Myerson [16]) *There exists a dual vector α such that $(\alpha_i * s_i)(t_i) > 0$ iff t_i jeopardizes s_i.*

and

Lemma 4.6 (Nau & McCardle [18]; Myerson [16]) *If a pure strategy profile s is coherent, then $f(s, \alpha) = 0$ for every dual vector α. If a pure strategy profile s is incoherent, then there exists a dual vector α such that $f(s, \alpha) > 0$.*

Finally, the set of dual vectors of a game is bounded and defined by a set of linear inequalities, hence it is a polytope. A dual vector is *interior* if it belongs to the relative interior of this polytope. Since the relative interior of a nonempty convex set is always non-empty, there always exists an interior dual vector. Furthermore:

Lemma 4.7 *If α is an interior dual vector then, for every player i and all pure strategies s_i and t_i of player i, $(\alpha_i * s_i)(t_i) > 0$ iff t_i jeopardizes s_i*

Proof. If α is an interior dual vector, then it satisfies with strict inequality all linear inequality constraints that are satisfied by all dual vectors and satisfied with strict inequality by at least one dual vector. This being noted, the result follows from the definition of a dual vector and lemma 4.5. ■

We are now in a position to prove proposition 4.1:

Proof of proposition 4.1. If a game is tight, then it follows from lemmas 4.3 and 4.7 that any interior dual vector satisfies the desired property. Conversely, if there exists a dual vector α such that, for all i and all s_i, $\alpha_i * s_i$ is completely mixed, then it follows from lemmas 4.5 and 4.3 that the game is tight. ■

4.3 Proof of proposition 4.2

We first need to introduce elements of dual reduction (Myerson [16]). *Throughout, α denotes a dual vector.* For every mixed strategy σ_i in $\Delta(S_i)$, define the mixed strategy $\alpha_i * \sigma_i$ by:

$$\forall t_i \in S_i, (\alpha_i * \sigma_i)(t_i) := \sum_{s_i \in S_i} \sigma_i(s_i) \left[(\alpha_i * s_i)(t_i)\right]$$

The transition probability α_i induces a Markov chain on S_i. This Markov chain partitions S_i into a set of transient states and disjoint recurrent classes:

$$S_i = T_i \coprod \left(\coprod_{1 \leq k \leq K} R_{i,k} \right)$$

where $\coprod$ denotes disjoint union, T_i is the (possibly empty) set of transient states, K a positive integer, and $R_{i,k}$ a recurrent class. A mixed strategy σ_i in $\Delta(S_i)$ is α_i-*invariant* if $\alpha_i * \sigma_i = \sigma_i$. For each recurrent class $R_{i,k} \subseteq S_i$, there exists a unique α_i-invariant mixed strategy with support in $R_{i,k}$, and its support is exactly $R_{i,k}$. Let S_i/α_i denote the set α_i-invariant mixed strategies with support in some recurrent class. It may be shown that a mixed strategy is α_i-invariant iff it is a convex combination of the strategies in S_i/α_i.

Definition (Myerson [16]) *The* α-reduced game G/α *is the game obtained from* G *by restricting player* i *to its* α_i*-invariant strategies. That is,*

$$G/\alpha = \{I, (S_i/\alpha_i)_{i\in I}, (U_i)_{i\in I}\}$$

Note that, since the pure strategies of G/α are mixed strategies of G, the mixed strategies of G/α may be seen as mixed strategies of G. As a particular case of (Myerson [16, theorem 1]), we have:

Lemma 4.8 *Any Nash equilibrium of* G/α *is a Nash equilibrium of* G.

Definition *A pure strategy* $s_i \in S_i$ *is* recurrent *under* α_i *if it belongs to a recurrent class of the Markov chain on* S_i *induced by* α_i*; that is, if there exists* σ_i *in* $S_i \backslash \alpha_i$ *such that* $\sigma_i(s_i) > 0$. *Otherwise* s_i *is* transient *under* α_i.

Lemma 4.9 *For every* i *in* I*, there exists a coherent pure strategy of player* i *which is recurrent under* α_i.

Proof. Let σ be a Nash equilibrium of G/α, hence also of G. Any pure strategy in the support of σ_i is both coherent (because σ is a Nash equilibrium, hence a correlated equilibrium of G) and recurrent under α_i (because $\sigma_i \in S_i/\alpha_i$). ■

Recall (4.1) and, for every mixed strategy profile σ, let

$$f(\sigma, \alpha) := \sum_{s\in S} \sigma(s) f(s, \alpha) \tag{4.2}$$

It follows from simple manipulations of the right-hand-side of (4.2) that (Myerson [16])

$$f(\sigma, \alpha) = \sum_{i\in I} [U_i(\alpha_i * \sigma_i, \sigma_{-i}) - U_i(\sigma)] \tag{4.3}$$

Define a mixed strategy profile σ to be α-invariant if $\alpha_i * \sigma_i = \sigma_i$ for every i in I. It is immediate from (4.3) that:

Lemma 4.10 ([18]; [16]) *If σ is α-invariant, then $f(\sigma, \alpha) = 0$.*

This allows to show that:

Lemma 4.11 *(i) Let $s \in S$. If $f(s, \alpha) > 0$ then for every σ in S/α we have $\sigma(s) = 0$.*

(ii) Let $s_i \in S_i$. If for every s_{-i} in S_{-i} we have $f(s, \alpha) > 0$, then for every σ_i in S_i/α_i we have $\sigma_i(s_i) = 0$; that is, s_i is transient under α_i.

(iii) For every i in I, let $S'_i \subseteq S_i$ and let $S' := \times_{i \in I} S_i$. Assume that for all s in $S \backslash S'$, $f(s, \alpha) > 0$. Then every s_i in $S_i \backslash S'_i$ is transient under α_i.

Proof. Proof of (i): Let $\sigma \in S/\alpha$. By definition of S/α, σ is α-invariant, hence by lemma 4.10, $f(\sigma, \alpha) = \sum_{s \in S} \sigma(s) f(s, \alpha) = 0$. Since $f(s, \alpha) \geq 0$ for all s in S by definition of a dual vector, this implies that $\sigma(s) = 0$ for every s in S such that $f(s, \alpha) > 0$.

Proof of (ii): Let $s_i \in S_i$. Assume that for every $s_{-i} \in S_{-i}$, $f(s, \alpha) > 0$. By (i), this implies that for every $s_{-i} \in S_{-i}$ and every σ in S/α, $\sigma(s) = 0$. This implies that $\sigma_i(s_i) = 0$ for every σ_i in S_i/α_i.

Proof of (iii): Let $s_i \in S_i \backslash S'_i$. The assumption implies that for every $s_{-i} \in S_{-i}$ we have $f(s, \alpha) > 0$. This being seen, (iii) follows from (ii). ∎

Finally:

Lemma 4.12 *If α is an interior dual vector then, for every incoherent pure strategy profile s in S, $f(s, \alpha) > 0$.*

Proof. As the proof of lemma 4.7, up to replacement of lemma 4.5 by lemma 4.6. ∎

We are now in a position to prove proposition 4.2:

Proof of proposition 4.2. Consider a pre-tight game. Let $S'_i = S^c_i$ and let α be an interior dual vector. It follows from lemma 4.12 that condition (B) is satisfied. We now prove that condition (A) is satisfied. By lemma 4.9, there exists a coherent pure strategy s_i which is recurrent under α_i. It follows from lemmas 4.4 and 4.7 that the support of $\alpha_i * S_i$ contains S_i^c. Furthermore, any pure strategy in the support of $\alpha_i * S_i$ is recurrent. But it follows from condition (B), lemma 4.11 item (iii), and from $S_i' = S_i^c$, that every pure strategy in $S_i \backslash S_i^c$ is transient. Therefore the support of $\alpha_i * S_i$ is exactly $S_i^c = S_i'$. Since this implies that any coherent pure strategy is recurrent, the same reasoning applies to any s_i in S_i'. This shows that condition (A) is satisfied.

We have shown that if a game is pre-tight then, for any interior dual vector and for $S_i' = S_i^c$, conditions (A) and (B) are satisfied. Conversely, assume that there exists a dual vector α and, for every player i in I, a subset S_i' of S_i such that conditions (A) and (B) are satisfied. Assume first that $S_i' = S_i^c$ for every i in I. In view of lemmas 4.4 and 4.5, condition (A) then implies that the game is pre-tight. Thus to prove that the game is pre-tight, it suffices to show that $S_i' = S_i^c$; this will also prove that if conditions (A) and (B) are satisfied then $S_i' = S_i^c$, as asserted in the last sentence of the proposition.

Let $s_i \in S_i \backslash S_i'$. For every s_{-i} in S_{-i}, $s = (s_i, s_{-i}) \notin \times_{i \in I} S_i'$. Therefore, $f(s, \alpha) > 0$ by condition (B). By lemma 4.6 (contraposition of the first sentence), this implies that the strategy profile s is incoherent. Since this holds for any strategy profile of the players other than i, it follows that the pure strategy s_i is incoherent. Therefore,

$$S_i^c \subseteq S_i' \tag{4.4}$$

It remains to prove the reverse inclusion. Condition (A) implies that S_i' is a recurrent class. Furthermore, it follows from condition (B) and lemma 4.11, item (iii), that the pure strategies of player i that do not belong to S_i' are transient under α_i. Therefore, S_i' is the unique recurrent class. This implies that there exists a unique α_i-invariant strategy σ_i and that its support is S_i'. In the reduced game G/α, the corresponding strategy profile $\sigma = (\sigma_i)_{i \in I}$ is the unique strategy profile, hence, trivially, a Nash equilibrium. By lemma 4.8, this implies that σ is a Nash equilibrium of G. Therefore, any pure strategy in the support of σ_i is coherent, i.e. $S_i' \subseteq S_i^c$. Together with (4.4), this shows that $S_i' = S_i^c$. This completes the proof. ■

Before turning to applications of propositions 4.1 and 4.2, we prove a result that will be used in later sections. It is more conveniently proved here as its proof is related to the proof of proposition 4.2.

4.4 Pre-tight games have quasi-strict equilibria

Definition A Nash equilibrium σ is *quasi-strict* if for every player i in I, any pure best-response to σ_{-i} belongs to the support of σ_i.

Games with 3 or more players need not have a quasi-strict Nash equilibrium (see van Damme [23, fig. 3.4.1 p. 56]). However:

Proposition 4.13 *(i) Any pre-tight game has a quasi-strict Nash equilibrium with support* $S^c := \times_{i\in I} S_i^c$.

(ii) Any tight game has a completely mixed Nash equilibrium

Proof. *Proof of (i):* Let G be a pre-tight game and let α be an interior dual vector. As shown in the proof of proposition 4.2 (first paragraph), conditions (A) and (B) of proposition 4.2 are satisfied. Therefore, if follows from the proof of proposition 4.2 (last paragraph) that in the reduced game G/α there is a unique strategy profile σ, that σ is a Nash equilibrium of G with support S^c and that if $s \notin S^c$, then $f(s,\alpha) > 0$. We now show that σ is quasi-strict: let s_i be a pure strategy of player i that does not belong to the support of σ_i. Since σ_i has support S_i^c, s_i is incoherent. This implies that for every s_{-i} in S_{-i}, $s = (s_i, s_{-i}) \notin S^c$, hence $f(s,\alpha) > 0$. Therefore, letting $\tau_i = s_i$ and $\tau_j = \sigma_j$ for $j \neq i$,

$$f(\tau,\alpha) = \sum_{s_{-i}\in S_{-i}} \tau(s) f(s,\alpha) > 0 \tag{4.5}$$

(The first equality merely recalls the definition of $f(\tau,\alpha)$). It follows from (4.3) and (4.5) that

$$\sum_{k\in I} [U_k(\alpha_k * \tau_k, \tau_{-k}) - U_k(\tau)] > 0 \tag{4.6}$$

Since for all $j \neq i$, $\tau_j = \sigma_j$ is α_j-invariant, and since $\tau_i = s_i$, (4.6) boils down to

$$U_i(\alpha_i * s_i, \sigma_{-i}) - U_i(s_i, \sigma_{-i}) > 0$$

Therefore s_i is not a best response to σ_{-i}. This shows that σ is quasi-strict.

Proof of (ii): Consider a tight game. By proposition 4.1, the game satisfies the conditions of proposition 4.2 with $S'_i = S_i$, hence $S_i = S^c_i$. Therefore it follows from (i) and from proposition 3.6 that the game has a completely mixed Nash equilibrium. ■

4.5 Examples

This section illustrates the use of propositions 4.1 and 4.2 and provides more examples of tight and pre-tight games.

Example 4.14 (A pre-tight game) Consider the following game, due to Bernheim [4] and studied by Nau and McCardle [18]:

$$\begin{array}{c} \\ T \\ M \\ B \end{array} \begin{array}{c} \begin{array}{ccc} L & M' & R \end{array} \\ \begin{pmatrix} 0,7 & 2,5 & 7,0 \\ 5,2 & 3,3 & 5,2 \\ 7,0 & 2,5 & 0,7 \end{pmatrix} \end{array}$$

Define α by $\alpha_1{*}T = \alpha_1{*}M = \alpha_1{*}B = M$ and $\alpha_2{*}L = \alpha_2{*}M' = \alpha_2{*}R = M'$. Let $S'_1 = \{M\}$ and $S'_2 = \{M'\}$. As noted, with another terminology, by Nau and McCardle [18, example 2]), α is a dual vector. Furthermore, if $s_1 \neq M$ and $s_2 \neq M'$, then $f(\alpha, s) = 3$. If $s_1 \neq M$ or $s_2 \neq M'$ (but not both), then $f(\alpha, s) = 1$. Thus, in any case, if $s \notin S'_1 \times S'_2$ (i.e. $s \neq (M, M')$), then $f(\alpha, s) > 0$. By proposition 4.1 this implies that the game is pre-tight. Proposition 4.1 also implies that $S^c_i = S'_i$ for $i = 1, 2$; that is, that (M, M') is the unique correlated equilibrium of the game, as noted by Nau and McCardle [18].

Example 4.15 (General Rock-Paper-Scissors games) A Rock-Paper-Scissors game is a 3×3 symmetric game in which the second strategy (Paper) beats the first (Rock), the third (Scissors) beats the second, and the first beats the third. Up to normalization (i.e. putting zeros on the diagonal) the payoff matrix of player 1 is of the form:

$$\begin{array}{c} \\ 1 \\ 2 \\ 3 \end{array} \begin{array}{c} \begin{array}{ccc} 1 & 2 & 3 \end{array} \\ \begin{pmatrix} 0 & -a_2 & b_3 \\ b_1 & 0 & -a_3 \\ -a_1 & b_2 & 0 \end{pmatrix} \end{array} \quad \text{with } a_k > 0,\ b_k > 0 \text{ for all } k = 1,2,3. \quad (4.7)$$

Note that we consider general Rock-Paper-Scissors games and not only the standard case $a_i = b_i = K$ for $i = 1, 2, 3$, where K is a positive constant.

Proposition 4.16 *Any Rock-Paper-Scissors game (4.7) is tight.*

Proof. Let {1,2,3} denote the set of pure strategies of both players (recall that the game is symmetric). Assume without loss of generality that $a_k + b_k < 1$ for all k in $\{1, 2, 3\}$. Counting k modulo 3, define the transition probability α_i as follows, for $i = 1, 2$: α_i maps the pure strategy k on the mixed strategy consisting in playing $k+1$ with probability a_k, $k-1$ with probability b_k and k with the remaining probability $1 - a_k - b_k$. Note that $\alpha_i * k$ is completely mixed, for every player i in $\{1, 2\}$ and every pure strategy k in $\{1, 2, 3\}$. Thus, by proposition 4.1, it suffices to check that $\alpha := (\alpha_1, \alpha_2)$ is a dual vector to prove that the game is tight. Due to the symmetry and cyclic symmetry of both α and the game, it is enough to check that $f(\alpha, s)$ is nonnegative for $s = (1, 1)$ and $s = (1, 2)$. For $s = (1, 1)$ we get:

$$f(\alpha, s) = 2(a_1[b_1] + b_1[-a_1]) = 0$$

For $s = (1, 2)$, we get

$$f(\alpha, s) = (a_1[a_2] + b_1[a_2 + b_2]) + (a_2[-b_1 - a_1] + b_2[-b_1]) = 0 \qquad \blacksquare$$

Example 4.17 (An n-player game) The following example (an n-player version of Matching Pennies) generalizes an example which appeared in an earlier version of (Nau et al. [17]). Consider an n-player game G_n in which every player has two pure strategies: K (for Keep) and R (for Reverse). The payoff of player $i \in \{1, 2, ..., n\}$ is $(-1)^{i+r}$ where r is the number of players playing R.

Proposition 4.18 *For every positive integer n, the game G_n is tight*

Proof. If n is even, define α by $(\alpha_i * R)(K) = (\alpha_i * K)(R) = 1/2$ for every i in $\{1, 2, ..., n\}$. If n is odd, hence $n = 2p + 1$, define α_i by $(\alpha_i * R)(K) = (\alpha_i * K)(R) = \frac{p+1}{2p+1}$ if i is even and by $(\alpha_i * R)(K) = (\alpha_i * K)(R) = \frac{p}{2p+1}$ if i is odd. It is easily checked that α is a dual vector. Furthermore, $\alpha_i * s_i$ is completely mixed for every player i in $\{1, 2, ..., n\}$ and every pure strategy s_i in $\{K, R\}$. By proposition 4.1, this implies that the game is tight. ■

5 Links between tight and pre-tight games

This section clarifies the links between tight and pre-tight games.

Proposition 5.1 *A game is tight iff it is pre-tight and every pure strategy of every player is coherent.*

Proof. If a game is pre-tight and if all pure strategies are coherent, then it follows from the definitions of tight and pre-tight games that the game is tight. Conversely, if a game is tight, then it is pre-tight, as noted in proposition 3.6, and it follows from proposition 4.13 that all pure strategies are coherent. ■

The next result shows that a game is pre-tight iff it becomes tight after deletion of all incoherent pure strategies. This motivates the choice of the term "pre-tight". We first need to introduce the game G^c obtained from G by restricting the players to their coherent strategies:

$$G^c = \{I, (S_i^c)_{i \in I}, (U_i)_{i \in I}\}$$

Proposition 5.2 *A game G is pre-tight iff the game G^c is tight.*

Proof. First, denote by $C^c \subseteq \Delta(S^c)$ the set of correlated equilibria of G^c. Since any correlated equilibrium of G has support in S^c, the set of correlated equilibria of G may be seen as a subset of $\Delta(S^c)$. Since the players have less possibilities of deviations in G^c than in G, it follows that any correlated equilibrium of G is a correlated equilibrium of G^c. That is, $C \subseteq C^c$.

Second, by definition 3.1, the game G^c is tight iff

$$\forall \mu \in C^c, h_{s_i,t_i}(\mu) = 0 \quad \forall i \in I, \forall s_i \in S_i^c, \forall t_i \in S_i^c \tag{5.1}$$

Similarly, by definition 3.5, G is pre-tight iff

$$\forall \mu \in C, h_{s_i,t_i}(\mu) = 0 \quad \forall i \in I, \forall s_i \in S_i^c, \forall t_i \in S_i^c \tag{5.2}$$

We need to show that (5.2) and (5.1) are equivalent. One sense is trivial: since $C \subseteq C^c$, it follows that (5.1) implies (5.2). We now show that (5.2) implies (5.1) by contraposition. Assume that (5.1) does not hold. Then:

$$\exists \mu \in C^c, \exists i \in I, \exists s_i^* \in S_i^c, \exists t_i^* \in S_i^c, h_{s_i,t_i}(\mu) > 0$$

Since $\mu \in C^c$, it follows that:

$$\forall i \in I, \forall s_i \in S_i, \forall t_i \in S_i^c, h_{s_i,t_i}(\mu) \geq 0 \tag{5.3}$$

(for $s_i \in S_i \backslash S_i^c$, this holds trivially as $\mu(s_i \times S_{-i}) = 0$). Furthermore, by lemma 4.13, there exists a quasi-strict Nash equilibrium σ with support S^c. Since σ is a Nash equilibrium, hence a correlated equilibrium,

$$\forall i \in I, \forall s_i \in S_i, \forall t_i \in S_i, h_{s_i,t_i}(\sigma) \geq 0 \tag{5.4}$$

Since σ is quasi-strict with support S^c,

$$\forall i \in I, \forall s_i \in S_i^c, \forall t_i \in S_i \backslash S_i^c, h_{s_i,t_i}(\sigma) > 0 \tag{5.5}$$

It follows from (5.3), (5.4) and (5.5) that, for $\epsilon > 0$ small enough, $\mu_\epsilon := \epsilon\mu + (1-\epsilon)\sigma$ is in C (to check that $h_{s_i,t_i}(\mu_\epsilon) \geq 0$, use (5.5) for t_i in $S_i \backslash S_i^c$ and (5.3) and (5.4) for t_i in S_i^c). But it follows from (5.4) and the definition of μ that $h_{s_i^*,t_i^*}(\mu_\epsilon) > 0$. This contradicts (5.2). ■

The reason why this implication (5.2) ⇒ (5.1) in the above proof is not trivial is that eliminating incoherent strategies can create new correlated equilibria. This occurs in the following example:

	x_2	y_2
x_1	$1, 1$	$0, 1$
y_1	$0, 1$	$1, 0$

	x_2	y_2
x_1	$1, 1$	$0, 1$

Let G denote the left game. It may be seen that G^c is the game on the right. In G^c any pure of mixed strategy profile is a Nash equilibrium. In G, a mixed strategy profile σ is a Nash equilibrium iff $\sigma_1(y_1) = 0$ and $\sigma_2(y_2) \leq 1/2$. Thus G^c has more Nash equilibria than G, hence also more correlated equilibria.

6 Topology of tight and pre-tight games

In this section we first show that the set of tight (resp. pre-tight) games is neither closed nor open; we then study the size of the class of tight (resp. pre-tight) games, i.e. whether it has Lebesgue measure 0 or not.

Example 6.1 *Consider the following* 2×2 *games:*

$$\begin{array}{cc} & \begin{array}{cc} L & R \end{array} \\ \begin{array}{c} T \\ B \end{array} & \begin{pmatrix} \epsilon, -\epsilon & 0,0 \\ 0,0 & \epsilon, -1 \end{pmatrix} \end{array} \qquad \begin{array}{cc} & \begin{array}{cc} L & R \end{array} \\ \begin{array}{c} T \\ B \end{array} & \begin{pmatrix} \epsilon, \epsilon & 0,0 \\ 0,0 & \epsilon, \epsilon \end{pmatrix} \end{array}$$

For $\epsilon > 0$, the left game is tight (apply proposition 4.1 with α defined by: $(\alpha_1 * T)(B) = (\alpha_2 * R)(L) = (\alpha_2 * L)(R) = \epsilon/2$ and $(\alpha_1 * B)(T) = 1/2)$. However, for $\epsilon = 0$, the left game is not even pre-tight, as in the Nash equilibrium (B, L), player 2 has a strict incentive not to play R, even though R is clearly coherent. This shows that the set of tight (resp. pre-tight) games is not closed. Furthermore, the game on the right is tight for $\epsilon = 0$, but for $\epsilon > 0$ it is not even pre-tight. This shows that the set of tight (resp. pre-tight) games is not open.

Another issue is the size of the class of tight (resp. pre-tight) games. Fix a positive integer n:

Proposition 6.2 *(i) Within the set of* $n \times n$ *bimatrix games, the set of tight games contains an open set. (ii) If* $n \neq m$*, then within the set of* $n \times m$ *bimatrix games, the set of tight games has Lebesgue measure 0.*

Proof. Proof of (i): Nitzan [19] shows that the set of $n \times n$ bimatrix games with a unique correlated equilibrium and such that this correlated equilibrium is a completely mixed Nash equilibrium, is nonempty and open. It follows from proposition 3.7 and proposition 5.1 that such games are tight, hence the result.

Proof of (ii): By item (ii) of proposition 4.13, any tight game has a completely mixed Nash equilibrium. Since, for $n \neq m$, the set of $n \times m$ games with a completely mixed Nash equilibrium has Lebesgue measure 0 (von Stengel [25, discussion following theorem 2.10]), this implies (ii). ■

In contrast with point (ii) of proposition 6.2, for any number of players n and any positive integers m_1, m_2,...,m_n:

Proposition 6.3 *The set of n-player pre-tight games of size $m_1 \times m_2 \times ... \times m_n$ contains a nonempty, open subset of the set of all n-player games of size $m_1 \times m_2 \times ... \times m_n$.*

Proof. It is shown in (Viossat, 2005) that the set of n-player games of size $m_1 \times m_2 \times ... \times m_n$ with a unique correlated equilibrium is a nonempty, open subset of the set of all n-player games of size $m_1 \times m_2 \times ... \times m_n$. Since any game with a unique correlated equilibrium is pre-tight, the result follows. ■

Thus, at least for $n \times m$ bimatrix games with $n \neq m$, the set of pre-tight games is much bigger than the set of tight games. Note however that for any positive integers n and m, within the set of $n \times m$ bimatrix games, the set of pre-tight games which do not have a unique correlated equilibrium has Lebesgue measure 0. This will be shown in section 8.

7 The geometry of Nash equilibria and correlated equilibria

As mentioned in the introduction, Nau et al. [17] proved the following:

Proposition 7.1 *If there is a Nash equilibrium σ in the relative interior of C, then:*

(a) The Nash equilibrium σ assigns positive probability to every coherent strategy of every player; that is, σ has support $S^c := \times_{i \in I} S_i^c$.

(b) The game is pre-tight.

For completeness, we recall the proof:

Proof. If (a) is not checked, then σ satisfies with equality some nonnegativity constraint which is not satisfied with equality by all correlated equilibria, hence σ belongs to the relative boundary of C. If condition (a) is checked then every coherent strategy of player i is a best-response to σ_{-i}. It follows that σ satisfies with equality all incentive constraints of type $h_{s_i,t_i}(\sigma) \geq 0$, where s_i and t_i are coherent. If the game is not pre-tight, at least one of these constraints is not satisfied with equality by all correlated equilibria, hence σ belongs to the relative boundary of C. ■

This section proves a converse of this result:

Proposition 7.2 *If a game is pre-tight, then C contains a Nash equilibrium in its relative interior.*

Proof. By proposition 4.13, there exists a quasi-strict Nash equilibrium σ with support S^c. This Nash equilibrium satisfies

$$\forall s \in S^c, \sigma(s) > 0 \tag{7.1}$$

and, by quasi-strictness,

$$\forall i \in I, \forall s_i \in S_i^c, \forall t_i \in S_i \backslash S_i^c, h_{s_i,t_i}(\sigma) > 0 \tag{7.2}$$

Since the inequalities in (7.1) and (7.2) are strict, there exists an neighbourhood Ω of σ in $\mathbb{R}^S$ in which (7.1) and (7.2) are still satisfied. Let E denote the linear subspace of $\mathbb{R}^S$ consisting of all vectors $x = (x(s))_{s \in S}$ such that

$$\sum_{s \in S} x(s) = 1 \text{ and } \forall s \in S \backslash S^c, x(s) = 0, \tag{7.3}$$

$$\forall i \in I, \forall s_i \in S_i \backslash S_i^c, \forall t_i \in S_i, h_{s_i,t_i}(x) = 0, \tag{7.4}$$

and

$$\forall i \in I, \forall s_i \in S_i^c, \forall t_i \in S_i^c, h_{s_i,t_i}(x) = 0. \tag{7.5}$$

Any correlated equilibrium satisfies trivially (7.3) and (7.4). Moreover, since the game is pre-tight, any correlated equilibrium satisfies (7.5). It follows that C is a subset of E. Furthermore, any vector in $\mathbb{R}^S$ satisfying the five conditions (7.1) to (7.5) is a correlated equilibrium. Therefore, $\Omega \cap E \subseteq C$. Since Ω is an open set containing σ and E a linear subspace containing C, it follows that σ belongs to the relative interior of C. ■

As an immediate consequence of propositions 7.1 and 7.2, we get:

Theorem 7.3 *The correlated equilibrium polytope of a game contains a Nash equilibrium in its relative interior iff the game is pre-tight.*

To conclude this section, note for later purposes that:

Proposition 7.4 *In a pre-tight game, a Nash equilibrium belongs to the relative interior of C iff it is quasi-strict.*

Proof. Let σ be a Nash equilibrium of a pre-tight game. By proposition 4.13, there exists a Nash equilibrium satisfying (7.1) and (7.2). If σ does not have support S^c, then it does not satisfy (7.1). If σ has support S^c but is not quasi-strict then it does not satisfy (7.2). In any case, if σ is not quasi-strict, then it satisfies with equality a nonnegativity or an incentive constraints that is satisfied with strict inequality by some correlated equilibrium. Therefore, σ does not belong to the relative interior of C.

Conversely, let σ be a quasi-strict Nash equilibrium. If σ has support S^c, then it follows from the proof of proposition 7.2 that σ belongs to the relative interior of C. But if $s_i \in S_i^c$ then, by definition of pre-tight games, s_i is a best-response to σ_{-i}. Since σ is quasi-strict, this implies that $\sigma_i(s_i) > 0$. It follows that σ has support S^c, completing the proof. ■

8 Two-player pre-tight games

In this section we first show that two-player zero-sum games are pre-tight but that a pre-tight game need not be best-response equivalent to a zero-sum game. We then show that, nevertheless, some of the properties of the equilibria and equilibrium payoffs of zero-sum games extend to pre-tight games.

8.1 Pre-tight games and zero-sum games

Proposition 8.1 *A two-player game which is best-response equivalent to a zero-sum game is pre-tight.*

Proof. In view of proposition 3.10 we only need to prove that two-player zero-sum games are pre-tight. Consider a two-player zero-sum game with value v. Let $\mu \in C$ and $s_i \in S_i$. As noted by Forges [7]:

(i) If $\mu(s_1 \times S_2) > 0$, then $\mu(\cdot|s_1)$ is an optimal strategy of player 2.

It follows that:

(ii) If a pure strategy of player 1 is coherent, then it is a best response to any optimal strategy of player 2.

Indeed, if t_1 is coherent then there exists μ in C and s_2 in S_2 such that $\mu(t_1|s_2)$ is positive. Assume that there exists an optimal strategy σ_2 of player 2 to which t_1 is not a best response. By playing σ_2 against $\mu(\cdot|s_2)$, player 2 would obtain strictly more than $-v$. Therefore $\mu(\cdot|s_2)$ is not an optimal strategy of player 1. This contradicts the analogue of (i) for player 2.

It follows from (i) and (ii) that in every correlated equilibrium μ and for every pure strategy s_1 played in μ, every coherent pure strategy of player 1 is a best-response to $\mu(\cdot|s_1)$. Together with the analogous result for player 2, this implies that the game is pre-tight. ■

The converse of proposition 8.1 is false:

Proposition 8.2 *A two-player tight game need not be best-response equivalent to a zero-sum game.*

Proof. Recall that every Rock-Paper-Scissors game is tight (proposition 4.16). We now show that Rock-Paper-Scissors games need not be best-response equivalent to a zero-sum game: In all bimatrix games that are best-response equivalent to a zero-sum game, fictitious play and its continuous time analog: the best-response dynamics, converge to the set of Nash equilibria (Robinson [20], Hofbauer and Sorin [13]). But, in Rock-Paper-Scissors games (4.7) such that $a_1a_2a_3 > b_1b_2b_3$, the best-response dynamics does not converge to the unique Nash equilibrium but to a triangle (see, for instance, Hofbauer and Sigmund [12]). The result follows. ■

The next section shows that, nevertheless, some of the main properties of two-player zero-sum games extend to pre-tight games. Noticeably, in two-player pre-tight games, the Nash equilibria are exchangeable and any correlated equilibrium payoff is a Nash equilibrium payoff.

8.2 Equilibria of pre-tight games

Let us first introduce some notations: we denote by NE the set of Nash equilibria of G and by NE_i the set of Nash equilibrium strategies of player

i. That is,

$$NE_i = \{\sigma_i \in \Delta(S_i), \exists \sigma_{-i} \in \times_{j \in I \setminus \{i\}} \Delta(S_j), (\sigma_i, \sigma_{-i}) \in NE\}$$

Proposition 8.3 *In a two-player pre-tight game:*

(a) NE_1 *and* NE_2 *are convex polytopes.*

(b) $NE = NE_1 \times NE_2$. *That is, Nash equilibria are exchangeable.*

We first need a lemma:

Lemma 8.4 *Let* G *be a two-player pre-tight game and let* $\sigma_1 \in \Delta(S_1)$. *The following assertions are equivalent:*

(i) σ_1 *is a Nash equilibrium strategy. That is,* $\sigma_1 \in NE_1$.

(ii) For some pure strategy s_2 *of player 2,* σ_1 *is the conditional strategy of player 1 given* s_2 *in some correlated equilibrium. Formally,* $\exists \mu \in C, \exists s_2 \in S_2, \mu(s_2 \times S_1) > 0$ *and* $\sigma_1 = \mu(\cdot|s_2)$.

(iii) Every pure strategy in the support of σ_1 *is coherent and all coherent pure strategies of player* 2 *are best responses to* σ_1.

(The analogous results for σ_2 *in* $\Delta(S_2)$ *hold obviously just as well.)*

Proof. (i) trivially implies (ii) and (ii) implies (iii) by definition of pre-tight games (definition 3.5). So we only need to prove that (iii) implies (i). Let σ_1 check (iii) and let $\tau_2 \in NE_2$. Necessarily, any pure strategy played in τ_2 is coherent. Since any coherent strategy of player 2 is a best response to σ_1, it follows that τ_2 is a best response to σ_1. Similarly, by the analogue of $(i) \Rightarrow (iii)$ for player 2, any coherent strategy of player 1 is a best response to τ_2. Since all pure strategies played in σ_1 are coherent, σ_1 is a best response to τ_2. Grouping these results, we get that (σ_1, τ_2) is a Nash equilibrium, hence $\sigma_1 \in NE_1$. ■

We now prove proposition 8.3: it follows from the proof of lemma 8.4 that if $\sigma_1 \in NE_1$, then for any $\tau_2 \in NE_2$, (σ_1, τ_2) is a Nash equilibrium. This implies that Nash equilibria are exchangeable (point (b)). Furthermore, from the equivalence of (i) and (iii) it follows that NE_1 can be defined by a finite number of linear inequalities. Therefore, NE_1 is a polytope, and so is NE_2 by symmetry (point (a)). ■

Our second result is that if μ is a correlated equilibrium, then the profile of its marginals is a Nash equilibrium:

Proposition 8.5 *Let μ be a correlated equilibrium of a two-player pre-tight game. Let $\sigma_i \in \Delta(S_i)$ denote the marginal probability distribution of μ on S_i. That is, $\forall s_i \in S_i, \sigma_i(s_i) = \mu(s_i \times S_{-i})$. Let $\sigma = (\sigma_1, \sigma_2)$ so that σ is the profile of the marginals of μ. We have:*

(a) σ is a Nash equilibrium

(b) The average payoff of the players is the same in σ and in μ. That is, $\forall i \in \{1, 2\}, U_i(\sigma) = U_i(\mu)$.

Proof. First note that σ_2 may be written:

$$\sigma_2 = \sum_{s_1 \in S_1 \,:\, \mu(s_1 \times S_2) > 0} \mu(s_1 \times S_2)\mu(\cdot|s_1) \tag{8.1}$$

Proof of (a): it follows from lemma 8.4 that for every $s_1 \in S_1$ with $\mu(s_1 \times S_2) > 0$, $\mu(\cdot|s_1) \in NE_2$. Therefore, by (8.1) and convexity of NE_2, $\sigma_2 \in NE_2$. Similarly, $\sigma_1 \in NE_1$, so that, by proposition 8.3, $\sigma \in NE$.

Proof of (b): assume $\mu(s_1 \times S_2) > 0$; then s_1 is coherent and, by definition of pre-tight games, any coherent strategy of player 1 is a best response to $\mu(\cdot|s_1)$. Since σ is a Nash equilibrium, every pure strategy in the support of σ_1 is coherent, so that

$$U_1(\sigma_1, \mu(\cdot|s_1)) = \max_{t_i \in S_i} U_1(t_1, \mu(\cdot|s_1)) = U_1(s_1, \mu(\cdot|s_1)) \tag{8.2}$$

Using successively (8.1), (8.2) and a straightforward computation, we get

$$\begin{aligned} U_1(\sigma) &= \textstyle\sum_{s_1 \in S_1 : \mu(s_1 \times S_2) > 0} \mu(s_1 \times S_2) U_1(\sigma_1, \mu(\cdot|s_1)) \\ &= \textstyle\sum_{s_1 \in S_1 : \mu(s_1 \times S_2) > 0} \mu(s_1 \times S_2) U_1(s_1, \mu(\cdot|s_1)) \;=\; U_1(\mu) \end{aligned}$$

Similarly, $U_2(\sigma) = U_2(\mu)$, completing the proof. ■

Finally, as noted by Forges [7], a two-player zero-sum game has a unique Nash equilibrium iff it has a unique correlated equilibrium. Since Bohnenblust et al. [5] showed that almost all zero-sum games have a unique Nash equilibrium, this implies that almost all zero-sum games have a unique correlated equilibrium. The next two propositions extend these results to two-player pre-tight games:

Proposition 8.6 *A two-player pre-tight game has a unique Nash equilibrium iff it has a unique correlated equilibrium.*

Proposition 8.7 *Within the set of $p \times q$ bimatrix games, the set of pre-tight games which do not have a unique correlated equilibrium has Lebesgue measure* 0*.*

Before proving these propositions note that, by proposition 6.3 the set of $p \times q$ pre-tight games contains a nonempty, open subset of the set of $p \times q$ bimatrix games. Therefore, proposition 8.7 implies that almost all pre-tight games have a unique correlated equilibrium. Note also, as an example of application of proposition 8.6, that since Rock-Paper-Scissors games 4.7 are tight (proposition 4.16) and have a unique Nash equilibrium (Hofbauer and Sigmund [12]), they have a unique correlated equilibrium.

Proof of propositions 8.6 and 8.7 If C is a singleton, then G has trivially a unique Nash equilibrium. Conversely, let G be a two-player pre-tight game such that C is not a singleton. By proposition 7.2, there exists a Nash equilibrium σ in the relative interior of C. Let τ be an extreme Nash equilibrium (in the sense of Evangelista and Raghavan [6]). Since, in two-player games, an extreme Nash equilibrium is an extreme point of C (Evangelista and Raghavan [6]), it follows that τ is an extreme point of C. Therefore $\tau \neq \sigma$. This proves proposition 8.6.[1] Furthermore, since τ does not belong to the relative interior of C, it follows from proposition 7.4 that τ is not quasi-strict. Since

[1] Proposition 8.6 also follows, and more directly, from lemma 8.4; but the above argument is convenient to prove jointly propositions 8.6 and 8.7.

almost all games have only quasi-strict equilibria (Harsanyi [10]), this implies proposition 8.7 ■

Proposition 8.8 *For almost all bimatrix games, either C is a singleton or all Nash equilibria belong to the relative boundary of C.*

Proof. This follows from theorem 7.3 and proposition 8.7. ■

The author does not know whether this result extends to games with three or more players. The reason why the proof for the two-player case does not go through is that, in games with three or more players, there need not be a Nash equilibrium that is an extreme point of C (Nau et al. [17]).

8.3 Equilibrium payoffs of pre-tight games

Let NEP (resp. NEP_i, CEP) denote the set of Nash equilibrium payoffs (resp. Nash equilibrium payoffs of player i, correlated equilibrium payoffs). That is,

$$NEP = \{g = (g_i)_{i \in I} \in \mathbb{R}^I \, : \, \exists \sigma \in NE, \forall i \in I, U_i(\sigma) = g_i\}$$

$$NEP_i = \{g_i \in \mathbb{R} \, : \, \exists \sigma \in NE, U_i(\sigma) = g_i\}$$

$$CEP = \{g = (g_i)_{i \in I} \in \mathbb{R}^I \, : \, \exists \mu \in C, \forall i \in I, U_i(\mu) = g_i\}$$

Two-player games which are best-response equivalent to zero-sum games may have an infinity of Nash equilibrium payoffs. So pre-tight games need not have a unique Nash equilibrium payoff. Nonetheless, some of the properties of equilibrium payoffs of zero-sum games are preserved. In particular, proposition 8.3 and proposition 8.5 imply respectively that:

Corollary 8.9 *In a two-player pre-tight game, NEP_1 and NEP_2 are convex and $NEP = NEP_1 \times NEP_2$*

Corollary 8.10 *In a two-player pre-tight game, $CEP = NEP$*

Thus, allowing for correlation is useless in two-player pre-tight games, in the sense that it cannot improve the equilibrium payoffs. In particular, there are no "good" correlated equilibria in the sense of Rosenthal [21]. Furthermore:

Proposition 8.11 *In a two-player pre-tight game, any correlated equilibrium payoff of player i given his move is a Nash equilibrium payoff of player i:*

$$\forall \mu \in C, \forall i \in \{1,2\}, \forall s_i \in S_i, \mu(s_i \times S_{-i}) > 0 \Rightarrow \sum_{s_{-i} \in S_{-i}} \mu(s_{-i}|s_i) U_i(s) \in NEP_i$$

Proof. Let $\mu \in C$ and $s_i \in S_i$ with $\mu(s_i \times S_{-i}) > 0$. Since $\mu \in C$, it follows that $U_i(s_i, \mu(\cdot|S_i)) = \max_{t_i \in S_i} U_i(t_i, \mu(\cdot|s_i))$. But by lemma 8.4, $\mu(\cdot|s_i) \in NE_i$. Therefore $\max_{t_i \in S_i} U_i(t_i, \mu(\cdot|s_i)) \in NEP_i$. The result follows. ■

Finally, there exists a dominant Nash equilibrium. That is,

Proposition 8.12 *There exists a Nash equilibrium σ such that*

$$\forall i \in \{1,2\}, U_i(\sigma) = \max NEP_i \tag{8.3}$$

Proof. Let τ, τ' be Nash equilibria such that $U_1(\tau) = \max NEP_1$ and $U_2(\tau') = \max NEP_2$. From exchangeability of equilibria, it follows that $\sigma = (\tau'_1, \tau_2)$ is a Nash equilibrium which satisfies (8.3). ■

8.4 Discussion

(a) Several classes of non-zero sum games in which some of the properties of two-player zero-sum games are satisfied have been studied. Most are defined in either of these three ways:

(i) by requiring some conflict in the preferences of the players over strategy profiles ("Strictly competitive games" (Aumann [1]; Friedman [8]), "Unilaterally competitive games" (Kats and Thisse [14]));

(ii) by comparing the best- or better-response correspondence in G and in some zero-sum game (games "order-equivalent" (Shapley [22]) or "best-response equivalent" (Rosenthal [21]) to a zero-sum game; "strategically zero-sum games" (Moulin and Vial [15]));

(iii) by comparing the Nash equilibria or Nash equilibrium payoffs of G and of some auxiliary game ("Almost strictly competitive games" (Aumann [1]) and other classes of games studied by Beaud [3]).

The definition of tight and pre-tight games do not fall in these categories; tight games however may be defined by comparing the *correlated equilibria* of G and of some auxiliary game. Indeed, let $-G$ be the game with the same sets of players and strategies as G but in which all the payoffs are multiplied by -1:

$$-G = \{I, (S_i)_{i\in I}, (-U_i)_{i\in I}\}$$

It is easily checked that G is tight iff G and $-G$ have the same correlated equilibria.

(b) Lemma 8.4 implies that in two-player *tight* games, as in two-player zero-sum games, the Nash equilibrium strategies of the players can be computed independently, as solutions of linear programs that depend only on the payoffs of the *other* player. In two-player *pre-tight* games, the additional knowledge of the sets of coherent strategies is required (indeed the 1×2 games (0 , 1 | 0 , 0) and (0 , 0 | 0 , 1) are both pre-tight and in both games the payoffs of player 1 are the same; but the Nash equilibrium strategies of player 2 are not the same).

(c) A wide range of dynamic procedures converge towards the set of correlated equilibria in all games (Hart [11]). By proposition 8.5, suitably modified versions of these dynamics converge towards the set of Nash equilibria in all two-player pre-tight games.

(d) In three-player tight games, Nash equilibria are not exchangeable. For instance, in the tight game from example 4.17, any mixed strategy profile in which two players randomize between their strategies with equal probability is a Nash equilibrium. Thus, for $n \geq 3$, if the Nash equilibria were exchangeable, then there would exist a pure Nash equilibrium. This is not the case.

Up to our knowledge, whether the other properties of section 8 extend to n-player games is not known.

Acknowledgments. This paper is based on chapter 6 of my PhD dissertation, written at the Laboratoire d'économétrie de l'Ecole polytechnique, Paris, and supervised by Sylvain Sorin. Part of the work was done at the Stockholm School of Economics. I am indebted to Galit Ashkenazi, Thomas Boulogne, Dinah Rosenberg, an associate editor, and especially to a student of Eilon Solan for highly relevant comments on previous versions of this paper. All errors are mine.

References

[1] Aumann R.J. (1961) "Almost Strictly Competitive Games", *Journal of the SIAM* **9**, 544-550

[2] Aumann, R.J. (1974) "Subjectivity and Correlation in Randomized Strategies", *Journal of Mathematical Economics* **1**, 67-96

[3] Beaud, J.P. (2002), "Antagonistic Games", *Mathématiques et Sciences Humaines*, **40**, 5-26

[4] Bernheim D.G. (1984), "Rationalizable Strategic Behavior", *Econometrica* **52**, 1007-1028

[5] Bohnenblust, H.F., S. Karlin and L.S. Shapley (1950), "Solutions of Discrete, Two-Person Games", in *Contribution to the Theory of Games*, H.W. Kuhn and A.W. Tucker Eds, Princeton University Press, vol. 1, 51-72

[6] Evangelista, F. and T.E.S. Raghavan (1996), "A Note on Correlated Equilibrium", *International Journal of Game Theory* **25**, 35-41

[7] Forges, F. (1990), "Correlated Equilibrium in Two-Person Zero-Sum Games", *Econometrica* **58**, 515

[8] Friedman J. (1983), "On Characterizing Equilibrium Points in Two-Person Strictly Competitive Games", *International Journal of Game Theory* **12**, 245-247.

[9] Gomez Canovas, S., P. Hansen, and B. Jaumard (1999), "Nash Equilibria from the Correlated Equilibria Viewpoint", *International Game Theory Review* **1**, 33-44

[10] Harsanyi, J.C. (1973), "Oddness of the Number of Equilibrium Points: A New Proof", *International Journal of Game Theory* **2**, 235-250

[11] Hart, S. (2005), "Adaptive Heuristics", *Econometrica* **73**, 1401-1430

[12] Hofbauer, J. and K. Sigmund (1998), Evolutionary Games and Population Dynamics, Cambridge University Press

[13] Hofbauer, J and S. Sorin (2006), "Best-Response Dynamics for Continuous Zero-Sum Games", Discrete and continuous dynamical systems-series B, **6**, 215-224

[14] Kats, A. and J.-F. Thisse (1992), "Unilateraly Competitive Games", *International Journal of Game Theory* **21**, 291-299

[15] Moulin, H and J-P. Vial (1978), "Strategically Zero-Sum Games: The Class of Games Whose Completely Mixed Equilibria Cannot be Improved Upon", *International Journal of Game Theory* **7**, 201-221

[16] Myerson, R.B. (1997), "Dual Reduction and Elementary Games", *Games and Economic Behavior* **21**, 183-202

[17] Nau R.F., S. Gomez Canovas, and P. Hansen (2004), "On the Geometry of Nash Equilibria and Correlated Equilibria", *International Journal of Game Theory* **32**, 443-453

[18] Nau, R.F. and K.F. McCardle (1990), "Coherent Behavior in Noncooperative Games", *Journal of Economic Theory* **50**, 424-444

[19] Nitzan, N. (2005), "Tight Correlated Equilibrium", discussion paper 394, Center for the Study of Rationality, the Hebrew University of Jerusalem

[20] Robinson, J. (1951). "An iterative method of solving a game", *Annals of Mathematics* **54**, 296-301

[21] Rosenthal, R.W. (1974), "Correlated Equilibria in Some Classes of Two-Person Games", *International Journal of Game Theory* **3**, 119-128

[22] Shapley, L. (1964), "Some Topics in Two-Person Games", *Annals of Mathematical Studies* **52**, Princeton, 1-28

[23] van Damme E. (1991), *Stability and Perfection of Nash Equilibria*, Springer-Verlag

[24] Viossat, Y (2005), "Openness of the Set of Games with a Unique Correlated Equilibrium", cahier du laboratoire d'économétrie 2005-28, Ecole polytechnique, France; revised in 2006

[25] von Stengel, B. (2002), "Computing Equilibria for Two-Person Games", in *Handbook of Game Theory*, R.J. Aumman and S. Hart Eds, Elsevier Science Publishers (North Holland), vol. 3, chap. 45, 1723-1759