Extreme Correlated and Nash Equilibria in Two-Person Games*

Martin W. Cripps

University of Warwick

Coventry CV4 7AL, UK

May 1992

This version May 2026

Abstract

I show that if σ is the correlated equilibrium distribution of a bimatrix game that is generated by an extreme Nash equilibrium, then σ is also an extreme point of the set of correlated equilibria of the bimatrix game. I describe a perturbation of a correlated equilibrium derived from an extreme Nash equilibrium that may generate further extreme points of the set of correlated equilibrium distributions. This can be used to generate simple conditions for the set of correlated equilibrium distributions to differ from the set of correlated equilibrium distributions that are derived from Nash equilibria.

Keywords: Correlated equilibrium; Nash equilibrium; extreme equilibrium; bimatrix games; zero-sum games.

JEL Classification Numbers: C70.

Address: Martin Cripps, Department of Economics, University of Warwick, Coventry CV4 7AL, UK.

*My thanks are due to Françoise Forges and to Jean-François Mertens for their encouragement, helpful suggestions, and for spotting an error in an earlier version of this paper; to Jonathan Cave and Ehud Kalai who also commented on this work; and to seminar participants at CORE. The existence of this improved version, May 2026, is entirely due to the encouragement and assistance of Fedor Sandomirskiy who wanted there to be a better online version of this unpublished paper.

1 Introduction

The set of Nash equilibrium strategy pairs for a bimatrix game is a finite union of convex polytopes, called maximal Nash subsets, and an extreme point a maximal Nash subset is called an *extreme equilibrium.* I will show that each extreme equilibrium generates a correlated equilibrium distribution that is an extreme point of the convex polytope of all correlated equilibrium distributions for the bimatrix game.

I will also provide a condition for the existence of a correlated equilibrium distribution that cannot be written as a convex combination of extreme equilibria and must therefore must involve private signals. Assume the game has a quasi-strict extreme equilibrium and restrict the players to use the pure strategies used at this equilibrium. Now choose a (mixed) strategy pair so that each player plays a worst response to their opponent's strategy on this sub-bimatrix (such a strategy pair is called an anti-Nash equilibrium here).[1] I will show that, provided the above two steps are possible and yield an anti-Nash equilibrium from the original restricted Nash equilibrium, then a linear combination of the correlated strategies generated by the anti-Nash equilibrium and the original extreme equilibrium generates a correlated equilibrium that is not a convex combination of the Nash equilibria.

2 The Derived Game

Define a finite bimatrix game where the row player has the pure strategies $I := \{1, 2, \ldots, m\}$, the column player has the pure strategies $J := \{1, 2, \ldots, n\}$, and the payoffs are described by the matrices (A, B), where $A := (a_{ij})_{i \in I,\, j \in J}$ and $B := (b_{ij})_{i \in I,\, j \in J}$. I will use (A, B) to denote this game. I will also use $p = (p_1, p_2, \ldots, p_m)^T \in \Delta^m$ and $q = (q_1, q_2, \ldots, q_n)^T \in \Delta^n$ to denote the players' mixed strategy vectors.[2] The set of all Nash equilibria of the game is

$$E(A,B) := \left\{ (p,q) \in \Delta^m \times \Delta^n : \begin{array}{ll} (p-p')^T A q \geq 0, & \forall p' \in \Delta^m \\ p^T B (q - q') \geq 0, & \forall q' \in \Delta^n \end{array} \right\}. \tag{1}$$

A correlated strategy for (A, B) induces a distribution over action pairs in $I \times J$. The distribution of a correlated strategy is denoted by $\sigma = (\sigma_{ij})_{i \in I, j \in J} \in \Delta^{mn}$. The set of all correlated equilibrium distributions of (A, B) is the set

$$C(A,B) := \left\{ \sigma \in \Delta^{mn} : \begin{array}{ll} \sum_j \sigma_{ij}(a_{ij} - a_{kj}) \geq 0, & \forall i, k \in I \\ \sum_i \sigma_{ij}(b_{ij} - b_{ih}) \geq 0, & \forall j, h \in J \end{array} \right\}. \tag{2}$$

When two players play the strategy pair (p, q), their play induces a probability distribution over action pairs. I define $\sigma(p, q)$ to be this distribution, that is, $\sigma_{ij}(p, q) := p_i q_j$ for all $i \in I$ and $j \in J$. Thus if $(p, q) \in E(A, B)$, then the correlated equilibrium derived from (p, q) is $\sigma(p, q)$ and the set of correlated equilibria that use only convex combinations of distributions

[1] My thanks are due to Jonathan Cave for suggesting this name to me.

[2] The superscript T denotes the transpose of a vector and $\Delta^k := \{x \in \mathbb{R}^k \mid x \geq 0,\ \mathbf{1}^T x = 1\}$, where $\mathbf{1} = (1, 1, \ldots, 1)^T$.

induced by Nash equilibria is

$$C_N(A,B) := \mathrm{co}\{\,\sigma(p,q) : (p,q) \in E(A,B)\,\} \subseteq \Delta^{mn}, \tag{3}$$

where $\mathrm{co}\{\cdot\}$ denotes the convex hull of the set. Aumann (1974) showed that $C_N(A,B) \subseteq C(A,B)$.

Hart and Schmeidler (1989) introduce an $mn \times (m^2+n^2)$ zero-sum game $(G,-G)$ with the payoffs derived from (A,B). The description of the payoffs in this game will be given later. We first describe the strategy sets of the players in $(G,-G)$ in more detail. The set of pure strategies for the row player in $(G,-G)$ is $I \times J$, therefore the row player's mixed strategies for $(G,-G)$ are correlated strategies for (A,B). I will also use $\sigma = (\sigma_{ij})_{i\in I,\, j\in J} \in \Delta^{mn}$ to describe the row player's mixed strategies in $(G,-G)$. A pure strategy for the column player in $(G,-G)$ is more complex. It first requires them to first choose a role: they can choose to be the row player in (A,B) or the column player in (A,B). If the column player in $(G,-G)$ choses to be the row player in (A,B), then they must select a pair of actions $(i,k) \in I^2$, and if they chose to be the column player then they must select a pair $(j,h) \in J^2$. The column's mixed strategies in $(G,-G)$ can therefore be described by the triples $(\rho,\pi,\tau) \in [0,1] \times \Delta^{mm} \times \Delta^{nn}$, where ρ is the probability that they choose to be the row player, $\pi := (\pi_{ik})_{i,k\in I} \in \Delta^{mm}$ is a probability measure on I^2, and $\tau := (\tau_{jh})_{j,h\in J} \in \Delta^{nn}$ is a probability measure on J^2. I will denote this mixed strategy by the vector θ using the convention $\theta^T = (\rho\pi^T, (1-\rho)\tau^T)$. In $(G,-G)$ the row player plays the mixed strategy σ over the rows of G and the column player uses θ over its columns.

We now describe the payoffs G in this derived game. Following Hart and Schmeidler (1989) the row player's expected payoffs in the zero-sum game $(G,-G)$ are the function

$$P(\sigma,\rho,\pi,\tau) = \rho \sum_{i,k\in I} \sum_{j\in J} \sigma_{ij}\pi_{ik}(a_{ij} - a_{kj}) + (1-\rho) \sum_{j,h\in J} \sum_{i\in I} \sigma_{ij}\tau_{jh}(b_{ij} - b_{ih}).$$

This payoff function can be written as a bilinear form, $P(\sigma,\rho,\pi,\tau) = \sigma^T G\theta$, where G is the partitioned matrix $G = [\,G_1\ G_2\,]$ given below. The submatrix G_1 is played when the column player in $(G,-G)$ chooses the role of row player in (A,B) and the submatrix G_2 is played when the column player in $(G,-G)$ chooses the role of column player in (A,B). The conventions used for the players' mixed strategies σ and θ give G_1 and G_2 their form.

$$G_1 = \begin{bmatrix} \begin{matrix} a_{11}-a_{11} & \cdots & a_{11}-a_{m1} \\ a_{12}-a_{12} & \cdots & a_{12}-a_{m2} \\ \vdots & & \vdots \\ a_{1n}-a_{1n} & \cdots & a_{1n}-a_{mn} \end{matrix} & & 0 \\ & \ddots & \\ 0 & & \begin{matrix} a_{m1}-a_{11} & \cdots & a_{m1}-a_{m1} \\ a_{m2}-a_{12} & \cdots & a_{m2}-a_{m2} \\ \vdots & & \vdots \\ a_{mn}-a_{1n} & \cdots & a_{mn}-a_{mn} \end{matrix} \end{bmatrix}$$

$$G_2 = \begin{bmatrix} b_{11}-b_{11} & \cdots & b_{11}-b_{1n} & 0 & \cdots & 0 & 0 & \cdots \\ 0 & \cdots & 0 & b_{12}-b_{11} & \cdots & b_{12}-b_{1n} & 0 & \cdots \\ \vdots & & \vdots & 0 & \cdots & 0 & b_{13}-b_{11} & \cdots \\ b_{21}-b_{21} & \cdots & b_{21}-b_{2n} & \vdots & & \vdots & & \\ 0 & \cdots & 0 & b_{22}-b_{21} & \cdots & b_{22}-b_{2n} & 0 & \cdots \\ \vdots & & \vdots & 0 & \cdots & 0 & b_{23}-b_{21} & \cdots \\ & \vdots & & & \vdots & & & \\ b_{m1}-b_{m1} & \cdots & b_{m1}-b_{mn} & 0 & \cdots & 0 & 0 & \cdots \\ 0 & \cdots & 0 & b_{m2}-b_{m1} & \cdots & b_{m2}-b_{mn} & 0 & \cdots \\ & \vdots & & & \vdots & & & \end{bmatrix}$$

The $mn \times m^2$ matrix G_1 is block-diagonal. There are m blocks along the diagonal. The ith diagonal block of G_1 has m columns that give expected payoffs $-\sum_j \sigma_{ij}(a_{ij} - a_{kj})$ $(k = 1, 2, \ldots, m)$ to the column player in $(G, -G)$. Up to a sign change, these payoffs are identical to the linear functions in the definition of a correlated equilibrium distribution (2). The $mn \times n^2$ matrix G_2 does not have a block-diagonal structure because of the ordering chosen for σ. In columns $(k-1)n+1, (k-1)n+2, \ldots, kn$ there are positive entries in the k^{th}, $k+n^{\text{th}}$, $k+2n^{\text{th}}$, $\ldots$ rows, so there is a permutation of the rows of G_2 which will give it a block-diagonal structure. Once this permutation is performed, the expected payoffs in the diagonal blocks are $-\sum_i \sigma_{ij}(b_{ij} - b_{ih})$ $(h = 1, 2, \ldots, n)$, which are a sign-changed version of the other linear functions in the definition of a correlated equilibrium distribution (2). The importance of the derived game $(G, -G)$ is that σ is a correlated equilibrium distribution of (A, B) if and only if it is an optimal strategy for the zero-sum game $(G, -G)$.

Result 1 (Hart and Schmeidler)**.** *$\sigma \in C(A, B)$ iff σ is an optimal strategy for the row player in $(G, -G)$. The value of $(G, -G)$ is zero.*

Remark 1. The strictly mixed distribution σ^ϵ is an ϵ-correlated equilibrium (Myerson, 1986, p. 139) for (A, B), if $\sigma^\epsilon_{ij} > \epsilon$ implies that the row player does not benefit by deviating from action i when it attaches probabilities proportional to $(\sigma^\epsilon_{i1}, \sigma^\epsilon_{i2}, \ldots, \sigma^\epsilon_{in})$ to the actions of the column player. (A symmetric incentive compatibility condition applies to the column player.) Further, the limit σ of a sequence of ϵ-correlated equilibria is an *acceptable correlated equilibrium.* But (σ, θ) is a perfect equilibrium of $(G, -G)$ if there is a sequence of strictly mixed strategies $(\sigma^\epsilon, \theta^\epsilon)$ such that $(\sigma^\epsilon, \theta^\epsilon) \to (\sigma, \theta)$ as $\epsilon \to 0$, where the strategy σ (respectively θ) is a best response to θ^ϵ (respectively σ^ϵ) for all ϵ (van Damme, 1987, p. 27). By comparing the two definitions above it is clear that an acceptable correlated equilibrium is a perfect equilibrium of $(G, -G)$. The characterization of a correlated equilibrium distribution as a Nash equilibrium of $(G, -G)$ thus suggests another way of describing refinements of correlated equilibria. It also provides a way of generating new refinements of the correlated equilibrium concept by applying a refinement of Nash equilibrium to the equilibria of $(G, -G)$.

Remark 2. It is possible to describe the set of subjective correlated equilibrium distributions of (A, B) as the set of optimal strategies in a two-person zero-sum game. In this new game the row player chooses two distributions: one for the row player in (A, B), the other for the column player in (A, B). The second player in the new game verifies the incentive compatibility conditions for each of these distributions. The set of a posteriori correlated equilibria is more difficult to characterize, but this would be very useful to do because of the relationship between a posteriori payoffs and rationalizable payoffs established by Brandenburger and Dekel (1987).

Remark 3. The Hart and Schmeidler (1989) game $(G, -G)$ can be defined for general n-person games, however, it is not clear how to generalize the later sections of this paper to the n-person case as there is no simple characterization of the extreme Nash equilibria of these games, see for example Chin et al. (1974).

3 Extreme Correlated Equilibrium Distributions in Two-Person Games

In this section I state the Shapley–Snow Theorem (1950) and use it, in combination with Result 1, to describe the extreme points of the convex set $C(A, B)$. I also describe a result due to Kuhn (1961) and Vorob'ev (1958) which characterizes the extreme equilibria in $E(A, B)$. Following this I establish the main result of the paper and then give some simple further results.

The first step is to find necessary and sufficient conditions for the distribution σ to be an extreme point of the convex polytope $C(A, B)$ (recall that an extreme point of a convex set is a point that cannot be written as a convex combination of other points in the set). Result 1 is essential here because it shows that the set of correlated equilibrium distributions of (A, B) is identical to the set of optimal strategies for the row player in $(G, -G)$. The Shapley–Snow Theorem (1950) provides a necessary and sufficient characterization of the extreme points in each player's set of optimal strategies in a zero-sum game. Therefore, combining these results tells us that the extreme points of $C(A, B)$ are characterized by the Shapley–Snow Theorem applied to the zero-sum game $(G, -G)$. The payoffs in G are derived from the payoffs in (A, B), so this gives a necessary and sufficient characterization of the set of extreme points of $C(A, B)$ in terms of the bimatrix (A, B).

The second step is to relate the set of equilibria $E(A, B)$ to the extreme points of $C(A, B)$. For a bimatrix game (A, B) the set of equilibria, $E(A, B)$, is the union of a finite number of closed convex sets. These are called either maximal Nash subsets or equilibrium components, and each of these sets is the Cartesian product of two convex polytopes of Nash equilibrium strategies (Jansen, 1981). Each equilibrium component is therefore completely described by its extreme points, which are called extreme equilibria. There is a characterization of the extreme equilibria of the bimatrix game (A, B). The result is due to Vorob'ev (1958) and Kuhn (1961) and can also be found in Parthasarathy and Raghavan (1971) or Jansen (1981); it provides a necessary condition for the extreme equilibria within each maximal Nash subset. To sum up, the result of Kuhn and Vorob'ev can be used to characterize a property of the extreme equilibria of $E(A, B)$, whereas the Shapley–Snow result will provide a necessary and sufficient condition for the extreme points of $C(A, B)$.

The equilibrium components that make up $E(A, B)$ will now be described in greater detail. Suppose that S is a subset of $E(A, B)$ with the property $(p, q), (p', q') \in S$ implies $(p, q'), (p', q) \in S$; such a set S is called a Nash subset. An *equilibrium component*, or *maximal Nash subset*, is a Nash subset S with the property that there is no Nash subset S' that properly contains S. Each maximal Nash subset is the Cartesian product of two convex polytopes (Jansen, 1981). Thus maximal Nash subsets have a finite number of extreme points. I will call an extreme point of a maximal Nash subset an *extreme equilibrium*. (Notice that these extreme points are defined relative to the equilibrium component and are not necessarily extreme points of the larger set $E(A, B)$.)

Definition 1. The pair $(p, q) \in E(A, B)$ is an extreme equilibrium if (p, q) is an extreme point of an equilibrium component.

The main result of this paper shows that if a Nash equilibrium (p, q) is an extreme equilibrium in $E(A, B)$ (and thus satisfies Kuhn–Vorob'ev) it must generate a correlated equilibrium distribution $\sigma(p, q)$ that satisfies the necessary and sufficient conditions for an extreme point of $C(A, B)$. Therefore, the extreme points of $C_N(A, B)$ are extreme points of $C(A, B)$.[3] If all of the extreme points of $C_N(A, B)$ are extreme points of $C(A, B)$, it follows that no extreme point of $C_N(A, B)$ can be written as a non-trivial convex combination of points in $C(A, B)$. The result is, of course, not sufficient to ensure that $C_N(A, B) = C(A, B)$, because this is only true if all the extreme points of $C(A, B)$ are extreme points of $C_N(A, B)$. In general there may be extreme points of the set of correlated equilibrium distributions that are not in $C_N(A, B)$. These properties are illustrated in the following example.

Example 1.

$$(A, B) = \begin{bmatrix} (4, 4) & (1, 5) \\ (5, 1) & (0, 0) \end{bmatrix}$$

The strategic form (A, B) has three distinct Nash equilibria: $((1, 0), (0, 1))$; $((0, 1), (1, 0))$; $((\frac{1}{2}, \frac{1}{2}), (\frac{1}{2}, \frac{1}{2}))$. The equilibrium components are all disjoint singletons, so every Nash equilibrium of this game is also an extreme equilibrium. These equilibria induce, respectively, the optimal strategies $(0, 1, 0, 0)$, $(0, 0, 1, 0)$, $(\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$ for the row player in $(G, -G)$. These are also correlated equilibrium distributions in $C(A, B)$. A typical correlated equilibrium distribution for this game can be written as the vector $(x, y, z, 1 - x - y - z) \in \Delta^4$. Vectors in Δ^4 can be shown inside a tetrahedron where the vertices represent the distributions $x = 1$, $y = 1$, $z = 1$, $x = y = z = 0$. For the above game I have plotted the sets $C(A, B)$ and $C_N(A, B)$ in such a tetrahedron in Figure 1 below. The red set $C(A, B)$ consists of the convex hull of $C_N(A, B)$ (in blue) region and its two remaining extreme points at $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}, 0)$ and $(0, \frac{1}{3}, \frac{1}{3}, \frac{1}{3})$.

Now I quote two central results on bimatrix games that will be used repeatedly below. Proofs of these results can be found in Parthasarathy and Raghavan (1971) or in the original sources. The notation $\mathbf{1} := (1, 1, \ldots, 1)^T$ is used and the dimension of this vector should be obvious from the context.

[3] Note: All the extreme points of $C_N(A, B)$ must be generated by extreme equilibria.

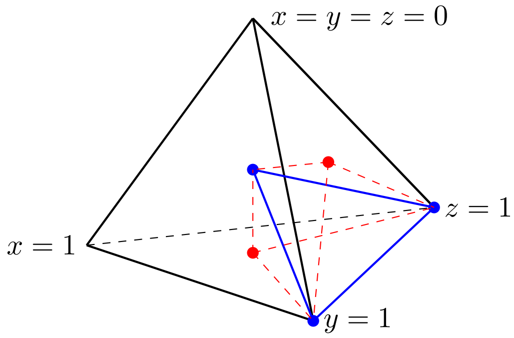

Figure 1: $C_N(A, B)$ and $C(A, B)$

Result 2 (Shapley–Snow)**.** *In a finite zero-sum game $(M, -M)$ with the value $v \neq 0$:*

(i) The set $E(M, -M)$ is convex and has a finite number of extreme points.

(ii) $(p, q) \in E(M, -M)$ is an extreme point of $E(M, -M)$ if and only if M has a non-singular submatrix M' such that the rows of M' contain the support of p, the columns of M' contain the support of q, and $M'q' = v\mathbf{1}$, $(p')^T M' = v\mathbf{1}^T$, where p' and q' are the restrictions of p and q to the rows and columns of M' respectively.

Result 3 (Kuhn–Vorob'ev)**.** *If (p, q) is an extreme equilibrium of (A, B), then there exist square submatrices A', B' of A and B such that:*

(i) The columns of A' are the support of q and the rows of B' are the support of p.

(ii) $A'q' = a\mathbf{1}$, $(p')^T B' = b\mathbf{1}^T$, where p', q' are vectors obtained from p, q with appropriate deletions of zeros and $a = p^T Aq$, $b = p^T Bq$.

(iii) The matrices H, K below are non-singular:

$$H = \begin{bmatrix} A' & -\mathbf{1} \\ \mathbf{1}^T & 0 \end{bmatrix}, \qquad K = \begin{bmatrix} B' & -\mathbf{1} \\ \mathbf{1}^T & 0 \end{bmatrix}.$$

The first proposition applies the Shapley–Snow result to $(G, -G)$ to characterize the extreme points of the set of correlated equilibrium distributions $C(A, B)$. The problem that this presents is that the value of the derived game $(G, -G)$ is zero and so Shapley–Snow cannot be directly applied. To avoid this problem I add a constant $\delta \neq 0$ to all the payoffs in G; let U be the matrix with unity in every entry, then $G + \delta U$ is the payoff matrix G with δ added to each element. This does not alter the set of equilibria, $E(G, -G) = E(G + \delta U, -G - \delta U)$, because the inequalities defining an equilibrium are linear in the payoffs. The value of $(G, -G)$ is zero, so the value of $(G + \delta U, -G - \delta U)$ will be δ.

Proposition 1. *σ is an extreme point of $C(A, B)$ if and only if σ is an optimal strategy for the row player in $(G, -G)$ and there exists a non-singular submatrix R of $G + \delta U$ with rows that contain the support of σ such that $(\sigma')^T R = \delta\mathbf{1}^T$, where σ' is the restriction of σ to the rows of R.*

Proof. If σ is an extreme point of $C(A,B)$, then it is an extreme point of the set of optimal strategies for the row player in $(G+\delta U, -G-\delta U)$ (by Result 1). The existence of R now follows from Result 2.

Let σ be an optimal strategy for the row player in $(G,-G)$ and $(\sigma')^T R = \delta \mathbf{1}^T$, where R is a non-singular submatrix of $G+\delta U$. If σ contains only one positive element, then it is a pure-strategy Nash equilibrium of (A,B) and thus an extreme point of $C(A,B)$. We henceforth assume that R has at least to columns.

(i) If R contains a column of δ's, then there exists an optimal strategy θ for the column player (play the column of δ's) such that $R\theta' = \delta\mathbf{1}$. By Result 2, (σ,θ) is an extreme point of $E(G+\delta U, -G-\delta U)$ and $E(G,-G)$. This set is a Cartesian product, so σ is an extreme point of the set of optimal strategies for the row player in $(G,-G)$, and by Result 1 σ is an extreme point of $C(A,B)$.

(ii) Suppose that R does not contain a column of δ's. As R is non-singular, there exists a unique vector x such that $Rx = \delta\mathbf{1}$. Suppose, without loss, that x has a non-zero element in its last entry. Define R' to be the matrix which replaces the last column of R with the vector $\delta\mathbf{1}$. (This is also a submatrix of $G+\delta U$ because $\delta\mathbf{1}$ is the first column of $G+\delta U$.) The matrix R' is also non-singular. To see this suppose it were not and that there exists $y \neq 0$ such that $R'y = 0$. Note that y must give non-zero weight to the final column, otherwise this implies that columns of R are linearly dependent. As the weight on the final column is non-zero $R'y = 0$ gives a non-zero combination of the columns of R (excluding the final one) that equals $\delta\mathbf{1}$. This is a contradiction as x is unique and gives positive weight to the final column.

This gives a non-singular submatrix R' of $G+\delta U$ and an optimal strategy θ (play the column of δ's) for the column player such that $(\sigma')^T R' = \delta\mathbf{1}^T$ and $R'\theta' = \delta\mathbf{1}$. By Result 2, (σ,θ) is an extreme point of $E(G+\delta U, -G-\delta U)$ and $E(G,-G)$. This set is a Cartesian product, so σ is an extreme point of the set of optimal strategies for the row player in $(G,-G)$, and by Result 1 σ is an extreme point of $C(A,B)$. ☐

Example 1 (Continued). In Figure 1, five extreme points are given for $C(A,B)$: $(0,1,0,0)$, $(0,0,1,0)$, $(\frac{1}{4},\frac{1}{4},\frac{1}{4},\frac{1}{4})$, $(\frac{1}{3},\frac{1}{3},\frac{1}{3},0)$, $(0,\frac{1}{3},\frac{1}{3},\frac{1}{3})$. Proposition 1 can be easily verified for the first two of these strategies because the relevant submatrices are 1-dimesional. A submatrix of $G+\delta U$ can also be found for the third, because the first, second, fourth, and sixth columns of $G+\delta U$ are linearly independent. The submatrices for the fourth and fifth strategies use columns one, two, six and columns one, four, seven respectively.

Proposition 2. *If $(p,q) \in E(A,B)$ is an extreme equilibrium, then $\sigma(p,q)$ is an extreme point of $C(A,B)$.*

Proof. Let $(p,q) \in E(A,B)$ be an extreme equilibrium (see Definition 1). Without loss of generality assume $p_i > 0$ iff $i = 1,2,\ldots,c \le m$ and $q_j > 0$ iff $j = 1,2,\ldots,d \le n$. Thus the distribution $\sigma(p,q)$ plays rows $hn+1, hn+2, \ldots, hn+d$ (for $h = 0,\ldots,c-1$) of the matrix $G+\delta U$ (where $\delta \neq 0$) with positive probability, that is cd rows in total. Suppose that $\sigma(p,q)$ is not extreme and so (by Proposition 1) it is not supported on a non-singular submatrix. Hence, there exists non-zero weights $\mu_1,\ldots,\mu_d,\mu_{n+1},\ldots,\mu_{n+d},\ldots,\mu_{(c-1)n+d}$ so that this linear combination of the cd rows of $G+\delta U$ sums to zero. As the first column of $G+\delta U$ is $\delta\mathbf{1}$ we have $\sum_t \mu_t = 0$ (where the summation ranges over all the weights μ_t).

Consider the rows $1, 2, \ldots, d$ of $G+\delta U$. These are the only rows in the support of $\sigma(p,q)$ with non-δ entries in the first m columns. Inspection of G_1 shows that the first m entries of the first d rows of $G+\delta U$ are the vectors $(\delta+a_{11})\mathbf{1}^T-(a_{.1})^T$, $(\delta+a_{12})\mathbf{1}^T-(a_{.2})^T$, $\ldots$, $(\delta+a_{1d})\mathbf{1}^T-(a_{.d})^T$, where $a_{.j}$ is the jth column of A. If the cd rows of $G+\delta U$ with weights μ_t sum to zero, then the sum over the first d rows must satisfy

$$\mu_1[(\delta+a_{11})\mathbf{1}-a_{.1}]+\mu_2[(\delta+a_{12})\mathbf{1}-a_{.2}]+\cdots+\mu_d[(\delta+a_{1d})\mathbf{1}-a_{.d}]+\sum_{t>d}\mu_t\delta\mathbf{1} \;=\; 0.$$

As $\sum_t \mu_t = 0$, this is equivalent to

$$\mu_1 a_{.1}+\mu_2 a_{.2}+\cdots+\mu_d a_{.d}-\left(\sum_{t\le d} a_{1t}\mu_t\right)\mathbf{1} \;=\; 0.$$

This is true if the weights $\mu_1, \mu_2, \ldots, \mu_d$ are all zero. Suppose instead that at least one of the weights $\mu_1, \ldots, \mu_d$ is non-zero. Recall that (p,q) is an extreme equilibrium and $(a_{.1}, a_{.2}, \ldots, a_{.d})$ are the columns that support q. They thus contain the submatrix A' described in Result 3. The above expression, therefore, gives a linear combination of the columns of the matrix H in Result 3 (omitting the final row) that sums to zero. (As well as some entries in the vectors these vectors that do not appear in H.) As H is non-singular there is no non-trivial combination of the columns of H that sum to zero, so the sum over the final row with these weights cannot be zero, or $\sum_{t\le d}\mu_t \neq 0$. Define $\lambda = (\lambda_1, \ldots, \lambda_d)^T$ to satisfy $\lambda_t = \mu_t/\sum_{t\le d}\mu_t$, then from above $A'\lambda = \alpha\mathbf{1}$, where $\alpha = \sum_{t\le d} a_{1t}\lambda_t$. The pair (p,q) is an extreme equilibrium, so by Result 3 $A'q' = a\mathbf{1}$, where q' is q with zero elements deleted. The vector $((q'-\lambda)^T, \alpha - a)$ therefore gives a linear combination of the columns of H so that they sum to zero. The columns of H are linearly independent, so this implies $q' = \lambda$, or $q_j = \mu_j/\sum_{t\le d}\mu_t$ for $j = 1, 2, \ldots, d$. Thus if one of the weights $\mu_1, \mu_2, \ldots, \mu_d$ is non-zero they must all have the same sign and be proportional to the equilibrium strategy q'.

Now repeat this argument for each of the c blocks of d rows in $G_1+\delta U$ that are in the support of $\sigma(p,q)$. This shows that for $i = 1, 2, \ldots, c$ the weights $(\mu_{in+1}, \mu_{in+2}, \ldots, \mu_{in+d})$ are either all zero or all have the same sign and be proportional to q'.

Now we link these c blocks of d rows together. There is a permutation of the rows which transforms G_2 into block-diagonal form. After this permutation, repeat the analysis above on the first block of c rows of $G_2+\delta U$. If $b_{i.}$ is the ith row of B, the weights $(\mu_1, \mu_{d+1}, \ldots, \mu_{(c-1)d+1})$ must satisfy

$$\mu_1 b_{1.}+\mu_{d+1}b_{2.}+\cdots+\mu_{(c-1)d+1}b_{c.}-\left(\sum_{i\le c} b_{i1}\mu_{(i-1)d+1}\right)\mathbf{1} \;=\; 0.$$

The matrix K of Result 3 is non-singular, so again it can be shown that either the vector $(\mu_1, \mu_{d+1}, \ldots, \mu_{(c-1)d+1})$ is zero or all of its elements have the same sign. Thus all the μ_t's in all of the c blocks of d rows either have the same sign or are zero. If one of the μ_t's is non-zero and all of the μ_t have the same sign, then it cannot be that $\sum_t \mu_t = 0$. Thus $\mu_t = 0$ for all t. The rows of $G + \delta U$ in the support of σ are therefore linearly independent.

This shows that the rows of $G+\delta$ in the support of σ are linearly independent. Thus there exists a non-singular cd-dimensional submatrix R of $G + \delta U$ that includes only these rows. To use Proposition 1 to show that $\sigma(p, q)$ is an extreme point of $C(A, B)$ we must also show that R can be chosen so that $(\sigma')^T R = \delta \mathbf{1}^T$ (where σ' is $\sigma(p, q)$ with appropriate deletions). We will choose R to be a submatrix of $G_1 + \delta U$. G_1 has a block-diagonal structure with m blocks. The argument above selected a $d \times d$ submatrix A' of the first block and the columns used in constructing the matrix A' above all give the Nash equilibrium payoff, by Result 3 (ii). By repeating this argument for each of the c blocks that contain the support of $\sigma(p, q)$ one can build up a $cd \times cd$ submatrix of $G_1 + \delta U$. This is a non-singular cd-dimensional submatrix R of $G + \delta U$ such that $(\sigma')^T R = \delta \mathbf{1}^T$ (where σ' is σ with appropriate deletions). By Proposition 1, $\sigma(p, q)$ is an extreme point of $C(A, B)$. □

This establishes the main result: the extreme points of the set of correlated equilibrium distributions generated by Nash equilibria, $C_N(A, B)$, are also extreme points of the set $C(A, B)$. Some consequences of the relationship between the extreme points of $C(A, B)$ and the extreme equilibria are the two propositions below. In Proposition 3 I give a lower bound on the number of constraints (2) that are binding at an extreme point of $C(A, B)$. Proposition 4 is a consequence of Result 1; it is a simple result on the uniqueness of correlated equilibrium distributions which follows from a well-known result on unique Nash equilibria. (A correlated equilibrium distribution σ is said to be of full support if $\sigma_{ij} > 0$ for all $i \in I$, $j \in J$.) The result gives a way of checking whether a game has both a unique correlated and a unique Nash equilibrium. To extend this result one could investigate when the existence of a unique Nash equilibrium of full support ensures the existence of a unique correlated equilibrium distribution. This is known to be true for zero-sum games (see Forges 1990), but is not clear for the non-zero-sum case.

Proposition 3. *If $\sigma \in \Delta^{mn}$ is an extreme point of $C(A, B)$, then at least $\sharp\{(i, j) : \sigma_{ij} > 0\} - 1 + m + n$ of the constraints in the definition (2) are binding at σ.*

Proof. If σ is an extreme point of $C(A, B)$, then there exists a non-singular submatrix R of $G + \delta U$ such that $(\sigma')^T R = \delta \mathbf{1}^T$, where R has dimension at least $\sharp\{(i, j) \mid \sigma_{ij} > 0\}$. Each column of $G + \delta U$ whose expected payoff under σ equals δ represents one of the constraints (2) that binds at σ. As at most one of the columns of R can be $\delta\mathbf{1}$, and there are $m + n$ constraints in (2) that are trivially satisfied, there will be at least $\sharp\{(i, j) \mid \sigma_{ij} > 0\} - 1 + m + n$ binding constraints. □

Proposition 4. *If all of the correlated equilibrium distributions of the bimatrix game (A, B) are of full support and if* $\operatorname{rank}(G) = mn - 1$*, then the game (A, B) has a unique correlated equilibrium distribution.*

Proof. Suppose that all of the optimal strategies for the row player in the game $(G, -G)$ are of full support. By Parthasarathy and Raghavan (1971, Theorem 3.1.5, p. 61), if all the optimal strategies for the row player in an $r \times t$ zero-sum game with rank $r - 1$ and a zero value are of full support, then there is a unique optimal strategy for the row player. Thus there is a unique optimal strategy for the row player in $(G, -G)$ and a unique correlated equilibrium distribution. $\square$

This also provides a sufficient conditions for a game to have a unique Nash equilibrium, as $C_N(A, B) \subseteq C(A, B)$.

Corollary 1. *If all the correlated equilibrium distributions of the game (A, B) are of full support and if* $\text{rank}(G) = mn - 1$*, then (A, B) has a unique Nash equilibrium.*

In this section I have established that the extreme equilibria in the set $E(A, B)$ generate correlated equilibrium distributions that are extreme points of the set $C(A, B)$, and derived some easy consequences of this. Now I investigate when the sets $C_N(A, B)$ and $C(A, B)$ are equal. That is, when are correlated equilibrium distributions only convex combinations of distributions induced by Nash equilibria? As a consequence I present some results on what causes the differences between the set of correlated equilibrium distributions and $C_N(A, B)$.

4 Correlated Equilibrium Distributions and $C_N(A, B)$

The previous section shows that the set of correlated equilibria derived from public correlation, $C_N(A, B)$, is a convex polytope with extreme points that are also extreme points of $C(A, B)$. In this section I give one explanation for why $C_N(A, B) \neq C(A, B)$. Hence, the canoncial correlated equilibria of some games may require private recommendations rather than public randomization over Nash equilibria. The approach I take is to describe a set of perturbations of $C_N(A, B)$ that can generate correlated equilibrium distributions that are not in $C_N(A, B)$. However, this is not a complete explanation of the differences between $C(A, B)$ and $C_N(A, B)$, as Remark 4 shows.

The perturbations of $C_N(A, B)$ described in this section take as a starting point a particular quasi-strict extreme equilibrium $(p, q) \in E(A, B)$. The correlated equilibrium distribution $\sigma(p, q) \in C_N(A, B)$ is perturbed by taking a linear combination of $\sigma(p, q)$ and one other point in Δ^{mn}. The precise form of the linear combination used is $(1 - \mu)^{-1}[\sigma(p, q) - \mu\sigma(u, v)]$. Thus the perturbation is determined by a pair of strategies $(u, v) \in \Delta^m \times \Delta^n$ and a weight $\mu \in (0, 1)$ small enough that $\mu\sigma_{ij}(u, v) \leq \sigma_{ij}(p, q)$ for all ij. This linear combination reduces the probability $\sigma(p, q)$ attaches to the set of action profiles played by $\sigma(u, v)$, whilst increasing the probability attached to the set of action profiles not played by $\sigma(u, v)$. Since (u, v) is chosen so that each player uses a worst response with the subgame, the perturbation shifts probability away from outcomes played by these worst-responses.

The direction of the perturbation of $\sigma(p, q)$ is determined by a particular pair of strategies $(u, v) \in \Delta^m \times \Delta^n$. I will now describe how this pair is determined. Let $(\hat{A}, \hat{B})$ be the game (A, B) where the rows and columns not used by (p, q) are deleted, so the bimatrix $(\hat{A}, \hat{B})$ contains only the action profiles (i, j) that are played with positive probability by $\sigma(p, q)$. Choose $(\hat{u}, \hat{v}) \in E(-\hat{A}, -\hat{B})$, that is $(\hat{u}, \hat{v})$ is a Nash equilibrium for this sub-bimatrix with the signs of the payoffs switched. I call the pair $(\hat{u}, \hat{v})$ an *anti-Nash equilibrium* of $(\hat{A}, \hat{B})$, because $(\hat{u}, \hat{v})$ is a pair of strategies for $(\hat{A}, \hat{B})$ where each player uses a worst response to their opponent's strategy

$$(\hat{u}, \hat{v}) \in E(-\hat{A}, -\hat{B}) \quad \Leftrightarrow \quad \hat{u}^T \hat{A} \hat{v} \leq u^T \hat{A} \hat{v}, \ \forall u; \quad \hat{u}^T \hat{B} \hat{v} \leq \hat{u}^T \hat{B} v, \ \forall v.$$

We can therefore interpret the pair $(\hat{u}, \hat{v})$ as a pair of "bad" strategies in the submatrix $(\hat{A}, \hat{B})$. Now the direction of the perturbation $\sigma(u, v)$ is given by the pair $(u, v) \in \Delta^m \times \Delta^n$ which is $(\hat{u}, \hat{v})$ with appropriate zeros added.

The effect of the linear combination $(1-\mu)^{-1}(\sigma(p, q) - \mu\sigma(u, v))$ is to reduce the weight put on the set of worst responses when the players use the Nash equilibrium (p, q). It is difficult to suggest a simple story behind the correlation devices that generate the distribution $(1-\mu)^{-1}(\sigma(p, q) - \mu\sigma(u, v))$ because the linear combination used is not convex. The correlated equilibrium distribution could, perhaps, be generated by a two-stage device: at the first stage a public signal is announced and this is used to coordinate the players' behaviour on one extreme Nash equilibrium; at the second stage the players receive private signals that implement the distribution $(1-\mu)^{-1}(\sigma(p, q) - \mu\sigma(u, v))$. Two examples below show how the perturbations work in practice.

Example 1 (Continued). There is one completely mixed Nash equilibrium $(p, q) = ((\frac{1}{2}, \frac{1}{2}), (\frac{1}{2}, \frac{1}{2}))$ for this game, which is also a quasi-strict and extreme equilibrium in $E(A, B)$. As (p, q) is completely mixed, $(-\hat{A}, -\hat{B}) = (-A, -B)$. The game $(-A, -B)$ has two pure strategy Nash equilibria $(u, v) = ((1, 0), (1, 0))$ and $(u', v') = ((0, 1), (0, 1))$, and these are also extreme equilibria in $E(-\hat{A}, -\hat{B})$. One can describe (u, v) and (u', v') as anti-Nash equilibria of (A, B) since both players use a worst response to their opponent. If we take $\mu = \frac{1}{4}$, the linear combinations $(1-\mu)^{-1}[\sigma(p, q) - \mu\sigma(u, v)]$ and $(1-\mu)^{-1}[\sigma(p, q) - \mu\sigma(u', v')]$ give the correlated equilibrium distributions $(0, \frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ and $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}, 0)$, which are the two extreme points of $C(A, B)$ that are not extreme points of $C_N(A, B)$ (see Figure 1). In this example, using the anti-Nash equilibria to perturb $(p, q) = ((\frac{1}{2}, \frac{1}{2}), (\frac{1}{2}, \frac{1}{2}))$ is sufficient to find all the extreme points of $C(A, B)$ that are not extreme points of $C_N(A, B)$.

Example 2. A game (A, B) in strategic form is given below, together with one correlated equilibrium distribution σ:

$$(A, B) = \begin{bmatrix} (0,0) & (1,-1) & (1,1) \\ (-1,1) & (0,0) & (3,1) \\ (1,1) & (1,3) & (0,0) \end{bmatrix} \qquad \sigma = \begin{bmatrix} 1/8 & 1/8 & 1/8 \\ 1/8 & 1/8 & 1/8 \\ 1/8 & 1/8 & 0 \end{bmatrix}$$

The symmetric, quasi-strict Nash equilibrium $p = q = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ of (A, B) generates a correlated equilibrium distribution $\sigma(p, q)$ which is an extreme point of $C(A, B)$ (by Proposition 2). The pair of strategies $u = (0, 0, 1)$, $v = (0, 0, 1)$ is an extreme equilibrium of $(-A, -B)$, and hence they are also an anti-Nash equilibrium of (A, B). The linear combination $(9/8)(\sigma(p, q) - (1/9)\sigma(u, v))$ gives the correlated equilibrium distribution σ above.

The perturbation used in the example gives a correlated equilibrium distribution outside the set $C_N(A, B)$ only if the anti-Nash equilibria $E(-\hat{A}, -\hat{B})$ are different from the original extreme quasi-strict equilibrium (p, q). For the above perturbation to be feasible we therefore require $(-\hat{A}, -\hat{B})$ to have at least two equilibria.

Remark 4. Forges (1985, p. 143) presents an example of a game (A, B) where an action profile (i, j) is played with positive probability at a correlated equilibrium distribution, but the game (A, B) has no Nash equilibrium that plays (i, j) with positive probability. The perturbations suggested in this section will therefore not find all the correlated equilibrium distributions of a game. This is because the class of perturbations here always has the property that the players play action profiles that are used with positive probability at some Nash equilibrium.

References

Robert J. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1:67–96, 1974.

Robert J. Aumann. Correlated equilibrium as an expression of Bayesian rationality. *Econometrica*, 55:1–18, 1987.

Adam Brandenburger and Eddie Dekel. Rationalizability and correlated equilibria. *Econometrica*, 55:1391–1402, 1987.

H. H. Chin, T. Parthasarathy, and T. E. S. Raghavan. Structure of equilibria in n-person cooperative games. *International Journal of Game Theory*, 3:1–19, 1974.

Françoise Forges. Correlated equilibria in a class of repeated games with incomplete information. *International Journal of Game Theory*, 14:129–150, 1985.

Françoise Forges. Correlated equilibrium in two-person zero-sum games. *Econometrica*, 58: 515, 1990.

Sergiu Hart and David Schmeidler. Existence of correlated equilibria. *Mathematics of Operations Research*, 14:18–25, 1989.

M. J. M. Jansen. Regularity and stability of equilibrium points of bimatrix games. *Mathematics of Operations Research*, 6:530–550, 1981.

Harold W. Kuhn. An algorithm for equilibrium points in bi-matrix games. *Proceedings of the National Academy of Sciences of the United States of America*, 47:1657–1662, 1961.

Roger B. Myerson. Acceptable and predominant correlated equilibria. *International Journal of Game Theory*, 15:133–154, 1986.

T. Parthasarathy and T. E. S. Raghavan. *Some Topics in Two-Person Games*. American Elsevier, New York, 1971.

Lloyd S. Shapley and R. N. Snow. Basic solutions of discrete games. In H. W. Kuhn and A. W. Tucker, editors, *Contributions to the Theory of Games*, volume 1, pages 27–35. Princeton University Press, Princeton, 1950.

Eric E. C. van Damme. *Stability and Perfection of Nash Equilibria*. Springer-Verlag, Berlin, 1987.

N. N. Vorob'ev. Equilibrium points in bi-matrix games. *Theory of Probability and Its Applications*, 3:297–309, 1958. English translation.