

ON THE GEOMETRY OF NASH EQUILIBRIA AND CORRELATED EQUILIBRIA

By Robert Nau*, Sabrina Gomez Canovas**, and Pierre Hansen**

*Fuqua School of Business
Duke University
Durham NC 27708-0120 USA
robert.nau@duke.edu

**Group for Research in Decision Analysis (GERAD)
École des Hautes Etudes Commerciales and École Polytechnique de Montréal
3000, Chemin de la Côte-Sainte-Catherine
Montréal, Québec
CANADA H3T 2A7

November 18, 2003

Forthcoming in *International Journal of Game Theory*

© Spring-Verlag 2004

Abstract: It is well known that the set of correlated equilibrium distributions of an n -player noncooperative game is a convex polytope that includes all the Nash equilibrium distributions. We demonstrate an elementary yet surprising result: the Nash equilibria all lie on the boundary of the polytope.

JEL Classification: C720.

We are grateful to Françoise Forges, Dan Arce, the editors, and several anonymous referees for helpful comments. This research was supported by the National Science Foundation under grant 98-09225 and by the Fuqua School of Business.

1. INTRODUCTION

It is a curiosity in the history of game theory that the study of correlated equilibria has lagged far behind the study of Nash equilibria. At the time that Nash (1951) formulated the concept of an equilibrium in independent strategies, the use of correlated strategies in noncooperative games was already under investigation¹ and there was keen interest in applications of the newly developed theory and methods of linear programming. Yet more than 20 years elapsed before Aumann (1974, 1987) proposed the concept of an equilibrium in correlated strategies and gave examples showing that correlated equilibria are sometimes more efficient and more intuitively reasonable than Nash equilibria. The set of correlated equilibrium distributions is a convex polytope, hence correlated equilibria can be easily found by linear programming methods, and extreme points of the set of correlated equilibria have rational coordinates when the payoff matrix is rational. By comparison, the set of Nash equilibrium distributions may be nonconvex or disconnected or consist only of points with irrational coordinates (see the examples of section 4 and 5); and solving for Nash equilibria in games with three or more players may require nonlinear optimization or the solution of systems of nonlinear equations. The mathematical simplicity of correlated equilibria suggests that their existence should be provable using only tools of linear algebra, rather than powerful fixed-point theorems, yet another 15 years elapsed before the first such elementary existence proofs were discovered (Hart and Schmeidler 1989, Nau and McCardle 1990). More recently, the comparative geometry of Nash and correlated equilibria has been explored further, and it has been found that

¹ The use of correlated mixed strategies in 2-player games was discussed by Raiffa (1951), who noted: “it is a useful concept since it serves to convexify certain regions [of expected payoffs] in the Euclidean plane.” (p. 8)

in 2-player (bimatrix) games, all extremal Nash equilibria are also extremal correlated equilibria (Cripps 1995, Evangelista and Raghavan 1996, Gomez Canovas et al. 1999), although this result does not hold with more than 2 players.²

The purpose of this paper is to point out—“prove” is perhaps too strong a word—a more elementary fact that so far apparently has gone unnoticed, but which, once it is pointed out, is the second most obvious fact about the geometrical relation between Nash and correlated equilibria: *the Nash equilibria all lie on the boundary of the correlated equilibrium polytope*. This means that if the polytope is of full dimension, the Nash equilibria lie on its relative boundary.

2. MAIN RESULT

Let G denote a finite noncooperative game, let n denote the number of players, let S_i denote the set of pure strategies of player i , where $|S_i| \geq 2$ for all i , let $S = S_1 \times \dots \times S_n$ denote the set of all joint strategies (outcomes of G), and let $N = |S|$ denote the number of outcomes. Let s_i denote a pure strategy of player i and let $s = (s_1, \dots, s_n) \in S$ denote a joint strategy of all players. Let $u_i(s)$ denote the payoff (utility) of player i when joint strategy s is played, and let $u_i(d_i, s_{-i})$ denote the payoff to player i when she chooses strategy $d_i \in S_i$ while the others adhere to s .

Definition: The game G is *non-trivial* if $u_i(s) \neq u_i(d_i, s_{-i})$ for some player i , some $s \in S$, and some $d_i \in S_i$.

² In a 2-player game, the set of Nash equilibrium distributions is a finite union of convex polytopes in the product space of marginal probability distributions on strategies of individual players (Jansen 1981). The result proved by Cripps, Raghavan and Evangelista, and Gomez Canovas et al. is that the extreme points of these polytopes correspond to extreme points of the correlated equilibrium polytope in the higher-dimensional space of joint probability distributions. Our examples in sections 4 and 6 show that in 3-player games it is possible that none of the extreme points of the correlated equilibrium polytope is a Nash equilibrium.

A correlated equilibrium distribution of G is a vector π in \mathfrak{R}^N satisfying the following linear constraints (Aumann 1987):

$$\pi(s) \geq 0 \quad \text{for all } s \in S \quad (1a)$$

$$\sum_{s \in S} \pi(s) = 1 \quad (1b)$$

$$\sum_{s_{-i} \in S_{-i}} \pi(s)(u_i(s) - u_i(d_i, s_{-i})) \geq 0 \quad \text{for all } i \text{ and for all } s_i, d_i \in S_i. \quad (1c)$$

The first two constraints (1ab) define an $N-1$ dimensional simplex, henceforth denoted as Π , consisting of all probability distributions on joint strategies.³ The remaining inequalities (1c) are incentive constraints with the following interpretation: consider π as a commonly-known probability distribution of “recommended” joint strategies generated by a possibly-correlated randomizing device, and suppose that each player is informed only of her own component of the recommended joint strategy. Then the constraints (1c) require that, conditional on knowing that her own recommended strategy is s_i , player i should have no incentive to defect to any other strategy d_i , assuming that the other players adhere to their own recommendations.

The set of all correlated equilibrium distributions determined by (1abc) is a convex polytope, henceforth denoted as C , which is a proper subset of Π if the game is non-trivial (because non-triviality entails that at least one of the incentive constraints is not satisfied everywhere in Π). The polytope C is of *full dimension* if it has dimension $N-1$, the same as Π . A correlated equilibrium distribution π is on the *boundary* of C if it lies in a face of C whose dimension is less than $N-1$, which in turn is true if and only if π lies on a supporting hyperplane

³ Although the distributions lie in an $N-1$ dimensional subspace, it is convenient to represent them as vectors in \mathfrak{R}^N in order to treat all strategies symmetrically.

of C whose normal vector is non-constant, i.e., linearly independent from the total-probability constraint (1b). If C is of less than full dimension, then all of its points are boundary and it has no interior, but if it is not a singleton it still has a *relative* interior and a relative boundary. We will return to this point in section 6.

The set I of all joint probability distributions that are independent between players is defined by a system of nonlinear constraints, viz.

$$I = \{\boldsymbol{\pi} \in \Pi: \pi(s) = \pi_1(s_1) \times \dots \times \pi_n(s_n) \quad \forall s \in S\},$$

where π_i denotes the marginal probability distribution on S_i induced by $\boldsymbol{\pi}$. I includes all the vertices of the simplex Π (which correspond to pure strategies and are trivially independent), as well as faces of the simplex on which only one player uses a mixed strategy and segments along which the mixed strategies of all the players but one are fixed, but everywhere else it is locally nonconvex in the sense that a strictly convex combination of two independent joint distributions in which two or more players have distinct marginal distributions is not independent. In a 2×2 game, Π is a 3-dimensional tetrahedron and I is a 2-dimensional saddle. (See Figure 1 below.) In larger games, the dimensionality of I may be many orders lower than that of Π : the former has dimension $|S_1| + \dots + |S_n| - n$, whereas the latter has dimension $|S_1| \times \dots \times |S_n| - 1$.

The set of Nash equilibria is the intersection of C and I , which is non-empty by virtue of Nash's (1951) existence proof. We are interested in the geometry of this intersection: where in C may the independent distributions lie? The answer is given by

Proposition 1: In any finite, non-trivial game, the Nash equilibria are on the boundary of the correlated equilibrium polytope. If the polytope is of full dimension, the Nash equilibria are on its relative boundary.

Proof: If a Nash equilibrium is not completely mixed, it assigns zero probability to one or more joint strategies, hence it satisfies at least one of the non-negativity constraints (1a) with equality. If it is completely mixed, a Nash equilibrium renders every player indifferent among all of her own strategies, hence it satisfies all of the incentive constraints (1c) with equality, at least one of which is non-trivial if the game is non-trivial. Hence every Nash equilibrium satisfies at least one non-negativity constraint or non-trivial incentive constraint with equality, and that constraint (together with (1b)) determines a face of C whose dimension is less than $N-1$.

QED

Thus, an independent distribution cannot be an interior point of C , if C has a non-empty interior.

3. A GENERIC EXAMPLE: BATTLE OF THE SEXES

The coordination game known as “battle of the sexes” (Raiffa 1951, Luce and Raiffa 1957) has the following payoff matrix:

	Left	Right
Top	3, 2	0, 0
Bottom	0, 0	2, 3

As is well known, this game has three Nash equilibria, two of which are in pure strategies. Its correlated equilibrium polytope has five vertices, two of which are non-Nash.⁴ The geometry of these solutions is shown in Figure 1: the probability simplex is a tetrahedron, the independence set is a saddle, and the correlated equilibrium polytope is a hexahedron (a triangular dipyramid) that touches the saddle at exactly three points: the Nash equilibria.

⁴ The pure Nash equilibria are TL and BR and the completely mixed Nash equilibrium is $(3/5 T, 2/5 B) \times (2/5 L, 3/5 R)$. The two non-Nash extremal correlated equilibria are $(2/7 TL, 3/7 TR, 2/7 BR)$ and $(3/8 TL, 1/4 BL, 3/8 BR)$.

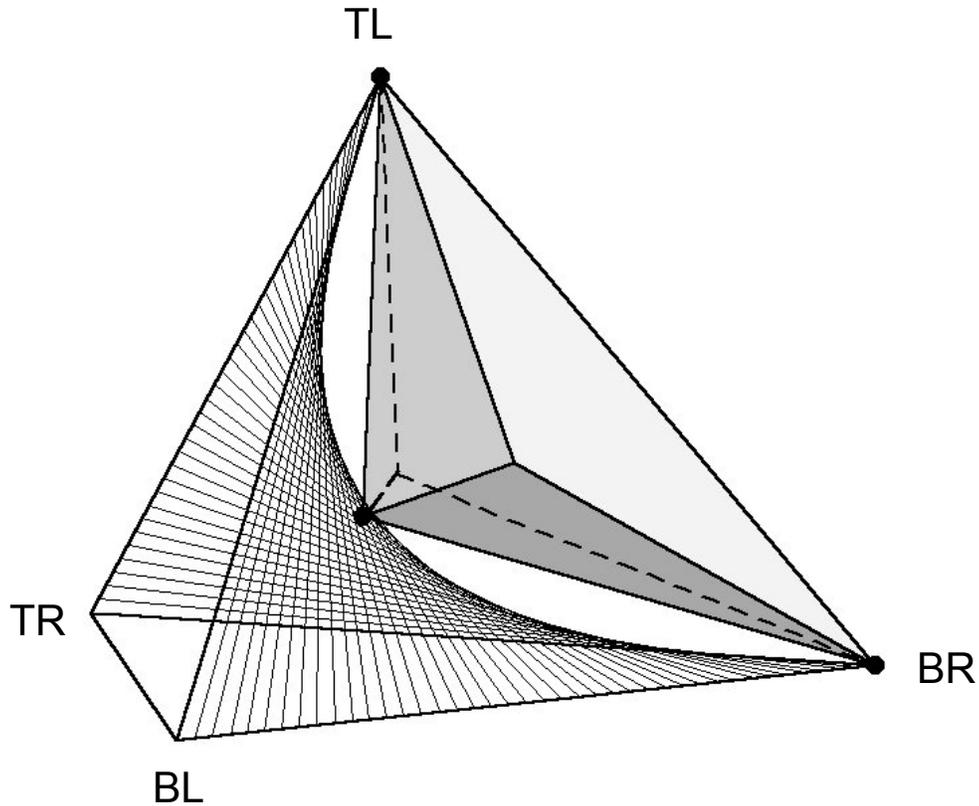


Figure 1. Geometry of the equilibria of “battle of the sexes”: the tetrahedron is the simplex of probability distributions on outcomes of the game, the saddle is the set of distributions independent between players, the polytope with 5 vertices and 6 facets is the set of correlated equilibria, and their three points of intersection are Nash equilibria.

This example is to some extent generic: a 2×2 game in which both players have distinct strategies that are not weakly dominated either has a correlated equilibrium polytope with five vertices, three of which are Nash equilibria arranged as in Figure 1, or else the polytope consists of a singleton, which may be either a pure-strategy or mixed-strategy Nash equilibrium. If some strategies lead to identical payoffs or are weakly dominated, then the polytope may have other numbers of vertices between one and five. For example, if the incentive constraints of one

player in battle-of-the-sexes are eliminated by equalizing the payoffs of her two strategies, the resulting correlated equilibrium polytope has four vertices.⁵

4. A THREE-PLAYER GAME WITH A UNIQUE NASH SOLUTION IN IRRATIONAL STRATEGIES

Nash (1951) gave an example of a 3-player poker game (devised by Lloyd Shapley) with a rational payoff matrix and a unique independent equilibrium in irrational mixed strategies. Such an equilibrium cannot be a vertex of the correlated equilibrium polytope, because the vertices must have rational coordinates, but according to Proposition 1 it still must lie somewhere on the boundary of the polytope. It is actually fairly easy to construct 3-player games with unique, irrational mixed-strategy Nash equilibria, such as:

	Left	Right
Top	3, 0, 2	0, 2, 0
Bottom	0, 1, 0	1, 0, 0

1

	Left	Right
Top	1, 0, 0	0, 1, 0
Bottom	0, 3, 0	2, 0, 3

2

(The numbers in the cells are the payoffs to Row, Column, and Matrix respectively.). The unique Nash equilibrium has the following marginal probabilities: $\pi(L) = (-13 + \sqrt{601})/24 \approx 0.480$, $\pi(T) = (9\pi(L) - 1)/(7\pi(L) + 2) \approx 0.619$, $\pi(1) = (-3\pi(L) + 2)/(\pi(L) + 1) \approx 0.379$.

⁵ While the 2x2 case is extremely simple to characterize, the number of vertices of the correlated equilibrium polytope may grow explosively with the size of the game. We randomly generated 2-player games of different sizes with non-negatively correlated payoffs and enumerated the vertices of their correlated equilibrium polytopes using Fukuda's (1993) implementation of the double-description method of Motzkin et al. (1953). Out of 250 4x4 games, half had polytopes with 5 or fewer vertices, but four games had polytopes with more than 100,000 vertices. However, the economically important vertices are those on the efficient frontier in expected payoff space, and their number grows much more slowly with the size of the game. In our sample of 4x4 games, more than three-quarters had a single efficient vertex (which happened also to be a Nash equilibrium), and the maximum number of efficient vertices was 65 (none of which was a Nash equilibrium).

The correlated equilibrium polytope of this game is seven-dimensional (i.e., full dimension) with 33 vertices. Even apart from the fact that the Nash equilibrium has irrational coordinates, it is clear that it cannot be a vertex of the polytope. The polytope is defined by a system of six incentive constraints (two for each player) in addition to the non-negativity and total probability constraints, and the incentive constraints are linearly independent. A completely mixed-strategy Nash equilibrium must satisfy all the incentive constraints with equality (because it renders every player indifferent among all her strategies), so it must be a point where all the incentive-constraint hyperplanes intersect. But, the probability simplex for this game is seven-dimensional, hence the intersection of the six hyperplanes only determines a line, rather than a point, in the linear span of the simplex. Since there exists a completely mixed Nash equilibrium, the set of correlated equilibrium distributions that satisfy all the incentive constraints with equality must be a line segment that passes through the interior of the probability simplex and terminates in two vertices on its boundary. The Nash equilibrium lies somewhere in the interior of this line segment, which is to say, it lies in the middle of an edge of the polytope.⁶ In particular, it is equal to $\alpha\pi^1 + (1-\alpha)\pi^2$, where $\alpha = 13((193\times\sqrt{601}) - (17\times 277))/(23\times 32) \approx 0.397$, and π^1 and π^2 are vertices of the polytope that assign the following probabilities to outcomes:

	TL1	TR1	BL1	BR1	TL2	TR2	BL2	BR2
π^1	0	24/78	9/78	0	24/78	0	5/78	16/78
π^2	216/1158	0	45/1158	144/1158	120/1158	384/1158	169/1158	80/1158

⁶ By the same reasoning, in any game where the number of incentive constraints is less than $N-1$, e.g., any game with 3 or more players in which each player has the same number of strategies (≥ 2), a completely mixed-strategy Nash equilibrium cannot be a vertex of the correlated equilibrium polytope.

5. A GAME WITH A CONTINUUM OF COMPLETELY MIXED-STRATEGY NASH EQUILIBRIA

The next example illustrates that not only can completely mixed-strategy Nash equilibria fall elsewhere than at vertices, but they can even form curves within faces of the polytope:

	Left	Right	
Top	0, 0, 2	0, 3, 0	
Bottom	3, 0, 0	0, 0, 0	
	1		

	Left	Right	
Top	1, 1, 0	0, 0, 0	
Bottom	0, 0, 0	0, 0, 3	
	2		

This game differs from the previous one in that the incentive constraints of the correlated equilibrium polytope are not all linearly independent: there are only five distinct incentive constraints, because the constraint for Row defecting from T to B is the same as the constraint for Column defecting from L to R. These five distinct constraints are independent, hence the set of points satisfying them with equality is two-dimensional. The correlated equilibrium polytope is seven-dimensional and has eight vertices. Three of the vertices (namely TR1, BL1, and BR2) are pure Nash equilibria, while two are incompletely mixed Nash equilibria: $(1/4 \text{ TR1}, 3/4 \text{ TR2})$ and $(1/4 \text{ BL1}, 3/4 \text{ BL2})$. The two incompletely mixed Nash equilibria satisfy all the incentive constraints with equality, as does the following extremal correlated equilibrium: $(3/20 \text{ TL1}, 1/10 \text{ BR1}, 9/20 \text{ TL2}, 6/20 \text{ BR2})$. The latter three vertices determine a face of the polytope that harbors a continuum of completely mixed Nash equilibria lying along an open curve, parameterized by $\pi(1)=1/4$ and $\pi(T)=(1-\pi(L))/(1-1/3\pi(L))$ for $0<\pi(L)<1$.

6. THE CASE OF LESS THAN FULL DIMENSION

The preceding three examples are *elementary* games in the sense of Myerson (1997). An elementary game has correlated equilibria that satisfy all the incentive and non-negativity constraints with strict inequality. The correlated equilibrium polytope of an elementary game has

full dimension, in which case Proposition 1 implies that the polytope has no Nash equilibria in its relative interior. As Myerson points out, elementary games are games for which it is unnecessary to consider refinements of correlated equilibrium, and every non-elementary game can be reduced to an elementary game by a process of *iterative dual reduction*, a generalization of the elimination of the weakly dominated strategies. If the correlated equilibrium polytope C is of less than full dimension, the game is not elementary, and the polytope has no interior, in which case Proposition 1 holds trivially. Nevertheless, if C is not a singleton, it has a *relative interior*, so the question remains: is it possible for a Nash equilibrium to exist in the relative interior of C when it is of less than full dimension? The answer is affirmative, but only for games that are rather special and seemingly uninteresting:

Proposition 2: A Nash equilibrium may exist in the relative interior of a correlated equilibrium polytope of less than full dimension only if the following conditions are satisfied:

- (i) The Nash equilibrium assigns positive probability to every coherent⁷ strategy of every player;
- (ii) In every correlated equilibrium, the incentive constraints for defecting from one coherent strategy to another coherent strategy are all satisfied with equality.

Proof: If the polytope is of less than full dimension, its relative boundary consists of the equilibria that satisfy with equality a non-negativity constraint (1a) or incentive constraint (1c) that is not satisfied with equality by all correlated equilibria. If condition (i) is not satisfied, then the Nash equilibrium satisfies with equality some *non-negativity* constraint that is not satisfied

⁷ An outcome of the game is defined to be *jointly coherent* if it does not lead to arbitrage when players reveal their utilities through side bets, which by linear duality is true if and only if that outcome has positive probability in some correlated equilibrium (Nau and McCardle 1990). A strategy of an individual player is therefore called *coherent* if it has positive probability in some correlated equilibrium.

with equality by all correlated equilibria, hence it is on the relative boundary. Assume, then, that condition (i) is satisfied by the Nash equilibrium and consider a player who has more than one coherent strategy. (There must be at least one such player, otherwise the polytope would be a singleton.) The Nash equilibrium strategy for that player is a mixture of the coherent strategies, and as such it must render the player indifferent among all those strategies—i.e., it must satisfy with equality all the incentive constraints for defecting from one coherent strategy to another coherent strategy. If condition (ii) is not also satisfied, then the Nash equilibrium satisfies with equality some *incentive* constraint that is not satisfied with equality by all correlated equilibria, hence it is on the relative boundary. *QED*

A game that satisfies the preceding conditions is one in which, in any equilibrium, every player is indifferent among all her coherent strategies given her recommended strategy, and the geometry of the set of correlated equilibria is highly non-robust to perturbations of the payoffs.

For example, consider the $2 \times 2 \times 4$ game:

	L	R		L	R		L	R		L	R
T	2,0,0	0,1,1	2	1,0,2	0,2,1	3	1,0,1	0,1,2	4	2,0,1	0,2,0
B	0,2,1	1,0,2	1	0,1,1	2,0,0	2	0,2,0	2,0,1	3	0,1,2	1,0,1
	1		2		3		4		5		6

The correlated equilibrium polytope is four-dimensional with six vertices. The set of Nash equilibria is a line segment terminating in two extremal Nash equilibria on the relative boundary of the polytope, and the interior of this line segment is in the relative interior of the polytope.

The extremal correlated and Nash equilibria have the following distributions over outcomes:

	TL1	TR1	BL1	BR1	TL2	TR2	BL2	BR2	TL3	TR3	BL3	BR3	TL4	TR4	BL4	BR4
Vertex #1	1/4			1/4		1/4	1/4									
Vertex #2		1/4	1/4		1/4			1/4								
Vertex #3									1/4			1/4		1/4	1/4	
Vertex #4										1/4	1/4		1/4			1/4
Vertex #5	1/8			1/8	1/8			1/8		1/8	1/8			1/8	1/8	
Vertex #6		1/8	1/8			1/8	1/8		1/8			1/8	1/8			1/8
Nash #1	1/8	1/8	1/8	1/8	1/8	1/8	1/8	1/8								
Nash #2									1/8	1/8	1/8	1/8	1/8	1/8	1/8	1/8

Nash equilibrium #1 is a convex combination of vertices #1 and #2, while Nash equilibrium #2 is a convex combination of vertices #3 and #4, hence the two extremal Nash equilibria lie on (disjoint) edges of the polytope but they are not extremal correlated equilibria. Any mixture of the two extremal Nash equilibria is a completely mixed Nash equilibrium in the relative interior of the polytope. The existence of the relative-interior Nash equilibria is a knife-edge situation, because a small perturbation of any one payoff destroys the symmetry⁸ of the game and causes the polytope to inflate to a higher-dimensional form that no longer satisfies condition (ii).

In the exceptional case where C has a Nash equilibrium in its relative interior, the set I of independent distributions still intersects it only “tangentially” in the following sense:

Proposition 3: Let G be a non-trivial game in which at least two players have two or more strategies. If the correlated equilibrium polytope C of G has a Nash equilibrium in its relative interior, then C lies on the relative boundary of a higher-dimensional correlated

⁸ Each 2×2 subgame induced by a pure strategy of the Matrix player has the same set of generic payoffs for each player, and subgames 2 and 4 are 180-degree rotations of subgames 1 and 3, respectively, while subgame 3 is obtained from subgame 2 by interchanging the payoffs of Row and Column and rotating 90 degrees clockwise.

equilibrium polytope C^* of a non-trivial game G^* having the same strategy space as G , where C^* has no Nash equilibria in its relative interior.

Proof: Consider two cases: (a) at least one player has an incoherent strategy in G , and (b) all strategies in G are coherent. In case (a), let G^* have the same strategy space as G , and let the payoffs in G^* of all players be equal to 0 in every outcome where all players use strategies that are coherent in G (regardless of whether the outcome is *jointly* coherent in G) and also in exactly one outcome—call it s^* —where all players except one use strategies that are coherent in G . Let the payoffs to all players be equal to -1 in the remaining outcomes where all players except one use strategies that are coherent in G (of which there is at least one); and let the payoffs to all players be equal to $-k$ in all outcomes (if any) where exactly k players use strategies that are incoherent in G . Then C^* is a simplex whose vertices assign probability 1 to the outcomes where the payoffs are equal to 0 (including s^*), and C is contained in a facet of C^* , namely the facet on which s^* has probability zero, which is part of the relative boundary of C^* . G^* has no Nash equilibria that assign positive probability to all its coherent strategies (because there is at least one outcome in which all players use coherent strategies yet which is jointly incoherent), hence by Proposition 2, part (i), the relative interior of C^* contains no independent distributions. Now consider case (b), in which all strategies of all players are coherent. Because C has less than full dimension, G has non-trivial incentive constraints, hence at least one player has two coherent strategies with distinct payoffs. If there is a Nash equilibrium in the relative interior of C , condition (i) of Proposition 2 requires it to be completely mixed, which implies that no coherent strategy weakly dominates another coherent strategy. Hence there is a player with two distinct coherent strategies that do not weakly dominate each other—say, strategies 1 and 2 of player 1. (There may be additional coherent strategies if $|S_1| > 2$.) It follows that $u_1(1, s_{-1}) -$

$u_1(2, s_{-1})$ has at least one strictly positive value and one strictly negative value as s_{-1} ranges over S_{-1} . Now let the game G^* with the same strategy space as G but with payoff functions $\{u_i^*\}$ be constructed as follows. For player 1, let $u_1^*(1, \cdot) = u_1(1, \cdot)$ and $u_1^*(k, \cdot) = u_1(2, \cdot)$ for $k \geq 2$. In other words, for player 1, strategies 1 and 2 have the same payoffs in G^* as in G , while all strategies $k > 2$ (if any) in G^* have the same payoffs as strategy 2 in G . For every other player $i > 1$, let $u_i^*(k, \cdot) = u_i(1, \cdot)$ for $k \geq 2$. In other words, for all players other than player 1, all strategies in G^* have the same payoffs as their strategy 1 in G . Then only player 1 has non-trivial incentive constraints in G^* . In particular, the incentive constraints of G^* are:

$$\sum_{s_{-1} \in S_{-1}} \pi(1, s_{-1})(u_1(1, s_{-1}) - u_1(2, s_{-1})) \geq 0, \quad (2a)$$

$$\sum_{s_{-1} \in S_{-1}} \pi(k, s_{-1})(u_1(2, s_{-1}) - u_1(1, s_{-1})) \geq 0 \quad \text{for } k = 2, \dots, |S_1|. \quad (2b)$$

These constraints are linearly independent (because the k^{th} constraint has non-zero coefficients only in outcomes where player 1 chooses strategy k), and each has at least one positive coefficient and one negative coefficient. The polytope C^* they determine has full dimension, so that (by Proposition 1) its relative interior excludes Nash equilibria. It is straightforward to show (invoking Proposition 2(ii)) that the elements of C satisfy (2ab) with equality, hence C lies in a proper face of C^* . *QED*

Thus, when C has a Nash equilibrium in its relative interior, the set I of independent distributions still touches C only from one “side,” namely from the outside of the higher-dimensional polytope C^* whose relative boundary contains C .

REFERENCES

- Aumann, R.J. (1974) "Subjectivity and Correlation in Randomized Games," *Econometrica* **30**, 445-462
- Aumann, R.J. (1987) "Correlated Equilibrium as an Expression of Bayesian Rationality," *Econometrica* **55**, 1-18
- Cripps, M. (1995) "Extreme Correlated and Nash Equilibria in Two-Person Games," working paper, University of Warwick
- Gomez Canovas, S., P. Hansen, and B. Jaumard (1999) "Nash Equilibria from the Correlated Equilibria Viewpoint," *International Game Theory Review* **1**, 33-44
- Evangelista, F. and T.E.S. Raghavan (1996) "A note on correlated equilibrium," *International Journal of Game Theory* **25**, 35-41.
- Hart, S. and D. Schmeidler (1989) "Existence of Correlated Equilibria," *Math. Oper. Res.* **14**, 18-25
- Jansen, M.J.M (1981) "Regularity and Stability of Equilibrium Points of Bimatrix Games," *Math. Oper. Res.* **6**, 18-25
- Fukuda, K. (1993). "cdd.c : C-implementation of the Double Description Method For Computing All Vertices And Extremal Rays Of A Convex Polyhedron Given By A System Of Linear Inequalities." Department of Mathematics, Swiss Federal Institute of Technology, Lausanne, Switzerland. The source code is currently available at http://www.cs.mcgill.ca/~fukuda/soft/cdd_home/cdd.html.
- Luce, R.D. and H. Raiffa (1957) *Games and Decisions*. New York: Wiley

- Motzkin, T.S., H. Raiffa, G.L. Thompson, and R.M. Thrall (1953) "The Double Description Method." In H.W. Kuhn and A.W. Tucker, editors, *Contributions to the Theory of Games*, Vol. 2. Princeton University Press
- Myerson, R.B. (1997) "Dual Reduction and Elementary Games." *Games and Economic Behavior* **21**, 183-202
- Nash, J. (1951) "Noncooperative Games." *Ann. Math.* **54**, 286-295
- Nau, R.F. and K. F. McCardle (1990) "Coherent Behavior in Noncooperative Games." *Journal of Economic Theory* **50**, 424-444
- Raiffa, H. (1951) "Arbitration Schemes For Generalized Two-Person Games." Ph.D. dissertation, University of Michigan