Strategic Characterization of the Index of an Equilibrium

Arndt von Schemde and Bernhard von Stengel

Department of Mathematics, London School of Economics, London WC2A 2AE, United Kingdom
schemde@gmail.com, stengel@nash.lse.ac.uk

Abstract. We prove that an equilibrium of a nondegenerate bimatrix game has index $+1$ if and only if it can be made the unique equilibrium of an extended game with additional strategies of one player. The main tool is the "dual construction". A simplicial polytope, dual to the common best-response polytope of one player, has its facets subdivided into best-response regions, so that equilibria are completely labeled points on the surface of that polytope. That surface has dimension $m-1$ for an $m \times n$ game, which is much lower than the dimension $m+n$ of the polytopes that are classically used.

1 Introduction

The index of a Nash equilibrium is an integer that is related to notions of "stability" of the equilibrium. In this paper, we only consider nondegenerate bimatrix games; "generic" (that is, almost all) bimatrix games are nondegenerate. A bimatrix game is nondegenerate if any mixed strategy with support of size k has at most k pure best responses [15]; the support of a mixed strategy is the set of pure strategies that are played positive probability. Nondegeneracy implies that the two strategies of a mixed equilibrium have supports of equal size. The index has the following elementary definition due to Shapley [13].

Definition 1. *Let (x,y) be a Nash equilibrium of a nondegenerate bimatrix game (A,B) with positive payoff matrices A and B, and let L and J be the respective supports of x and y, with corresponding submatrices A_{LJ} and B_{LJ} of the payoff matrices A and B. Then the* index *of (x,y) is defined as*

$$(-1)^{|L|+1}\text{sign}\ (\det(A_{LJ})\ \det(B_{LJ})). \tag{1}$$

The index has the following properties, which require that its sign alternates with the parity of the support size as in (1).

Proposition 2. *In a nondegenerate bimatrix game,*
(a) *the index of an equilibrium is $+1$ or -1;*
(b) *any pure-strategy equilibrium has index $+1$;*
(c) *the index only depends on the payoffs in the support of the equilibrium strategies;*
(d) *the index does not depend on the order of a player's pure strategies;*

B. Monien and U.-P. Schroeder (Eds.): SAGT 2008, LNCS 4997, pp. 242–254, 2008.
© Springer-Verlag Berlin Heidelberg 2008

(e) *the endpoints of any Lemke–Howson path have opposite index;*
(f) *the sum of the indices over all equilibria is* $+1$*;*
(g) *in a* 2×2 *game with two pure equilibria, the mixed equilibrium has index* -1.

Condition (a) holds because payoff-submatrices A_{LJ} or B_{LJ} that do not have full rank $|L|$ can only occur for degenerate games. The simple property (g) applies to, say, a coordination game and easily follows from (1) or (f). It is one indication that, as suggested by Hofbauer [7], equilibria of index $+1$ are in many respects "sustainable" according to Myerson [10], who discusses ways to refine or select equilibria in "culturally familiar games". Hofbauer [7] also shows that only equilibria of index $+1$ can be stable under any "Nash dynamics", that is, a vector field on the set of mixed strategy profiles whose rest points are the Nash equilibria [6][4]. Such dynamics may represent evolutionary or learning processes.

The most interesting computational property is (e), proved by Shapley [13]. The Lemke–Howson (LH) algorithm [9] (for an exposition see [15]) defines a path which can either start at a trivial "artificial equilibrium" with empty support, or else at any Nash equilibrium, and which ends at another equilibrium. The equilibria of the game, plus the artificial equilibrium, are therefore the endpoints of the LH paths. By (1), the artificial equilibrium has index -1. Consequently, the game has one more equilibrium of index $+1$ than of index -1, and (f) holds.

Equilibria as endpoints of LH paths provide a "parity argument" that puts the problem of finding one Nash equilibrium of a bimatrix game into the complexity class PPAD [11]. This stands for "polynomial parity argument with direction", where the direction of the path is provided by the index (which can also be determined locally at any point on the path).

The index of an equilibrium can also be defined for general games (which may be degenerate and have more than two players) as the degree of a topological map that has the Nash equilibria as fixed points, like the mentioned "Nash dynamics" [6][4].

The index is a relatively complicated topological notion, essentially a geometric orientation of the equilibrium. In this paper, we prove the following theorem, first conjectured in [7], which characterizes the index in much simpler strategic terms.

Theorem 3. *A Nash equilibrium of a nondegenerate* $m \times n$ *bimatrix game* G *has index* $+1$ *if and only if it is the unique equilibrium of a game* G' *obtained from* G *by adding suitable strategies. It suffices to add* $3m$ *strategies for the column player.*

The equilibrium of G in Theorem 3 is re-interpreted as an equilibrium of G', so none of the added strategies is used in the equilibrium; their purpose is to eliminate all other equilibria. Unplayed strategies do not matter for the index of an equilibrium by Prop. 2(c). By (f), a unique equilibrium has index $+1$, so only equilibria with positive index can be made unique as stated in Theorem 3; the nontrivial part is therefore to show that this is always possible.

We prove Theorem 3 using a novel geometric-combinatorial method that we call the *dual construction*. It allows to visualize all equilibria of an $m \times n$ game in a diagram of dimension $m-1$. For example, all equilibria of a $3 \times n$ game are visualized with a diagram (essentially, of suitably connected $n+3$ points) in the plane. This should provide new insights into the geometry of Nash equilibria.

A better understanding of the geometry of Nash equilibria may also be relevant algorithmically, and we think the index is relevant apart from providing the "D" in "PPAD". Recent results on the complexity of finding one Nash equilibrium of a bimatrix game have illustrated the difficulty of the problem: it is PPAD-complete [2], and LH paths may be exponentially long [12]. Even a sub-exponential algorithm for finding one Nash equilibrium is not in sight. In designing any such algorithm, for example incremental or by divide-and-conquer, it is important that the information carried to the next phase of the algorithm does not describe the entire set of equilibria, because questions about that set, for example uniqueness of the Nash equilibrium, tend to be NP-hard [5][3]. On the other hand, Nash equilibria with additional properties (for example, a minimum payoff) may not exist, or give rise to NP-complete decision problems. However, it is always possible to restrict the search to an equilibrium with index $+1$; whether this is of computational use remains speculative.

The dual construction has first been described in the first author's PhD dissertation, published in [14]. Some steps of the construction are greatly simplified here, and the constructive proof outlined in Section 5 is new.

2 Dualizing the First Player's Best Response Polytope

We use the following notation. All vectors are column vectors. If $d \in \mathbb{R}^k$ and $s \in \mathbb{R}$, then ds is the vector d scaled with s, as the product of a $k \times 1$ with a 1×1 matrix. If $s = 1/t$, we write d/t for ds. The vectors $\mathbf{0}$ and $\mathbf{1}$ have all components equal to 0 and 1, respectively. Inequalities like $d \geq \mathbf{0}$ between vectors hold for all components. A matrix C with all entries scaled by s is written as sC. We write $C = [c_1 \cdots c_k]$ if C is a matrix with columns $c_1, \ldots, c_k$. The transpose of C is $C^\top$.

The index of an equilibrium is defined in (1) via the sign of determinants. We recall some properties of determinants. Exchanging any two rows or any two columns of a square matrix $C = [c_1 \cdots c_k]$ changes the sign of $\det(C)$, which implies Prop. 2(d). The determinant is multilinear, so that, for any $d \in \mathbb{R}^k$, $s \in \mathbb{R}$ and $1 \leq i \leq k$,

$$\begin{aligned} \det[c_1 \cdots c_{i-1}\ c_i s\ c_{i+1} \cdots c_k] &= s \det(C), \\ \det[c_1 \cdots c_{i-1}\ (c_i + d)\ c_{i+1} \cdots c_k] &= \det(C) + \det[c_1 \cdots c_{i-1}\ d\ c_{i+1} \cdots c_k]. \end{aligned} \tag{2}$$

Let $s_1, \ldots, s_k$ be scalars, which we add to the columns of C. Repeated application of (2) gives

$$\det(C + [\mathbf{1}s_1 \cdots \mathbf{1}s_k]) = \det(C) + \sum_{i=1}^{k} s_i \det[c_1 \cdots c_{i-1}\ \mathbf{1}\ c_{i+1} \cdots c_k]. \tag{3}$$

The right-hand side of (3) is linear in $(s_1, \ldots, s_k)$. In particular, if $s_1 = \cdots s_k = s$, then the expression $\det(C + s\,[\mathbf{1} \cdots \mathbf{1}])$ is linear in s and changes its sign at most once.

We first explain why the matrices A and B are assumed to be positive in Def. 1. Consider an equilibrium (x,y), and discard for simplicity all pure strategies that are not in the support of x or y, so that A and B are the matrices called A_{LJ} and B_{LJ} in (1), which have full rank. Then the equilibrium payoffs to the two players are u and v, respectively, with $Ay = \mathbf{1}u$ and $B^\top x = \mathbf{1}v$. We want that always $u > 0$ and $v > 0$; this clearly holds if A and B have positive entries, although this not a necessary condition. Adding any constant t to all payoffs of A does not change the equilibria of the game, but does change the equilibrium payoff from u to $u+t$. Consequently, we could achieve $Ay = \mathbf{0}$ (with $t = -u$), or $Ay < \mathbf{0}$. However, $Ay = \mathbf{0}$ implies $\det(A) = 0$, and consequently a change of the sign of u implies a change of the sign of $\det(A)$. Because the sign of $\det(A)$ changes only once, that sign is unique whenever A is positive. Similarly, the sign of $\det(B)$ is unique, so (1) defines the index uniquely.

For the rest of the paper, we consider a nondegenerate $m \times n$ bimatrix game (A,B) so that the best-response payoff to any mixed strategy of a player is always positive, for example by assuming that A and B are positive. The following polytopes can be used to characterize the equilibria of (A,B) [15]:

$$\begin{aligned} P &= \{x \in \mathbb{R}^m \mid \quad x \geq \mathbf{0},\ B^\top x \leq \mathbf{1}\}, \\ Q &= \{y \in \mathbb{R}^n \mid Ay \leq \mathbf{1}, \qquad y \geq \mathbf{0}\}. \end{aligned} \tag{4}$$

Any $(x,y) \in P \times Q$ with $x \neq \mathbf{0}, y \neq \mathbf{0}$ represents a mixed strategy pair with best-response payoffs scaled to one; normalizing x and y as probability vectors re-scales these payoffs. The inequalities in P and Q are *labeled* with the numbers $1,\ldots,m+n$ to indicate the pure strategies of player 1 (labels $1,\ldots,m$) and player 2 (labels $m+1,\ldots,m+n$). Given $x \in P$ and $y \in Q$, each *binding* inequality (which holds as an equation) defines a facet of P or Q (by nondegeneracy, it cannot be a lower-dimensional face [15]). The corresponding label defines an unplayed own pure strategy or best response of the other player. An equilibrium of (A,B) corresponds a pair (x,y) of $P \times Q$ where each pure strategy $1,\ldots,m+n$ appears as a label of x or y. The artificial equilibrium is given by $(x,y) = (\mathbf{0},\mathbf{0})$.

The first step of our construction is to dualize the polytope P by considering its *polar* polytope [16]. Suppose R is a polytope defined by inequalities that has $\mathbf{0}$ it its interior, so that it can be written as $R = \{z \in \mathbb{R}^m \mid c_i^\top z \leq 1,\ 1 \leq i \leq k\}$. Then the polar polytope is defined as $R^\triangle = \operatorname{conv}\{c_i \mid 1 \leq i \leq k\}$, that is, as the convex hull of the normal vectors c_i of the inequalities that define R. The face lattice of $R^\triangle$ is that of R upside down, so $R^\triangle$ and R contain the same combinatorial information about the face incidences. More precisely, assuming that R has full dimension m, any face of R of dimension h given by $\{z \in R \mid c_i^\top z = 1 \text{ for } i \in K\}$ (with maximal set K) corresponds to the face $\operatorname{conv}\{c_i \mid i \in K\}$ of dimension $m-1-h$. So facets (irredundant inequalities) of R correspond to vertices of $R^\triangle$, and vertices of R correspond to facets of $R^\triangle$. If R is simple, that is, has no point that lies on more that m facets, then $R^\triangle$ is simplicial, that is, all its facets are simplices.

Because the game is nondegenerate, the polytope P is simple, and any binding inequality of P defines either a facet or the empty face, the latter corresponding to a dominated strategy of player 2 that can be omitted from the game. In particular, player 2 has no weakly dominated strategy, which would define a lower-dimensional face of P.

Because P does not have $\mathbf{0}$ in its interior, we consider instead the polytope

$$\begin{aligned} P_\varepsilon &= \{x \in \mathbb{R}^m \mid \qquad\qquad x \geq -\mathbf{1}\varepsilon,\ B^\top x \leq \mathbf{1}\}, \\ &= \{x \in \mathbb{R}^m \mid -1/\varepsilon \cdot Ix \leq \mathbf{1}, \qquad B^\top x \leq \mathbf{1}\}, \end{aligned} \tag{5}$$

where $\varepsilon > 0$ and I is the $m \times m$ identity matrix. For sufficiently small ε, the polytopes P and P_ε are combinatorially equivalent, because the simple polytope P allows small perturbations of its facets. Moreover, nondegeneracy crucially forbids weakly dominated strategies, which would be "better" than the dominating strategy under the "negative probabilities" x_i allowed in P_ε, and hence define facets of P_ε but not of P. Then

$$P_\varepsilon^\triangle = \mathrm{conv}(\{-e_i/\varepsilon \mid 1 \le i \le m\} \cup \{b_j \mid 1 \le j \le n\}), \tag{6}$$

where e_i is the ith unit vector in $\mathbb{R}^m$ and $B = [b_1 \cdots b_n]$. That is, $P_\varepsilon^\triangle$ is the convex hull of sufficiently largely scaled negative unit vectors and of the columns b_j of the payoff matrix B of player 2. We will later add points, which are just additional payoff columns; this is the reason why we perturb, rather than translate, P.

Any facet F of $P_\varepsilon^\triangle$ is a simplex, given as the convex hull of m vertices $-e_i/\varepsilon$ for $i \in K$ and b_j for $j \in J$, where $|K| + |J| = m$. We write

$$F = F(K,J) = \mathrm{conv}(\{-e_i/\varepsilon \mid i \in K\} \cup \{b_j \mid j \in J\}). \tag{7}$$

Then the vertices of the facet $F(K,J)$ define *labels* i and $m + j$ which represent unplayed strategies $i \in K$ of player 1 and best responses $j \in J$ of player 2. These labels are the labels of the facets of P_ε, and hence of P, that correspond to the vertices of $F(K,J)$.

The facet $F(K,J)$ itself corresponds to a vertex x_ε of P_ε. Namely, because $P_\varepsilon^{\triangle\triangle} = P_\varepsilon$ [16], we have $F(K,J) = P_\varepsilon \cap \{z \in \mathbb{R}^m \mid x_\varepsilon^\top z = 1\}$, where $x_\varepsilon^\top z \le 1$ holds for all $z \in P_\varepsilon$. The vertex x_ε of P_ε corresponds to a vertex x of P, which is determined from K and J by the linear equations $x_i = 0$ for $i \in K$ and $\sum_{i \notin K} b_{ij} x_i = 1$ for $j \in J$. The corresponding equations for x_ε are $(x_\varepsilon)_i = -\varepsilon$ for $i \in K$ and $\sum_{i=1}^m b_{ij}(x_\varepsilon)_i = 1$ for $j \in J$, so $x_\varepsilon \to x$ as $\varepsilon \to 0$.

In summary, the normal vectors of facets $F(K,J)$ of $P_\varepsilon^\triangle$ correspond to mixed strategies x of player 1. The vertices of such facets represent player 1's unplayed strategies $i \in K$ and best responses $j \in J$ of player 2. A similar representation of mixed strategies and best responses is considered by Bárány, Vempala and Vetta [1], namely the polyhedron defined as the intersection of the halfspaces with nonnegative normal vectors that contain the points $b_1, \ldots, b_n$. Our polytope $P_\varepsilon^\triangle$ approximates that polyhedron when it is intersected with the halfspace with supporting hyperplane through the m points $-e_i/\varepsilon$ for $1 \le i \le m$.

The facet $F_0 = F(\{1,\ldots,m\},\emptyset)$ of $P_\varepsilon^\triangle$ whose vertices are the m points $-e_i/\varepsilon$ for $1 \le i \le m$ corresponds to the vertex $\mathbf{0}$ of P. The surface of $P_\varepsilon^\triangle$ can be projected to F_0, giving a so-called Schlegel diagram [16]. A suitable projection point is $-\mathbf{1}/\varepsilon$. The Schlegel diagram is a subdivision of the simplex F_0 into simplices that correspond to the other facets of $P_\varepsilon^\triangle$. All these simplices have dimension $m-1$, so for $m = 3$ one obtains a subdivided triangle. An example is the left picture in Fig. 1 for the matrix B of the 3×4 game

$$A = \begin{bmatrix} 0 & 10 & 0 & 10 \\ 10 & 0 & 0 & 0 \\ 8 & 0 & 10 & 8 \end{bmatrix}, \qquad B = \begin{bmatrix} 0 & 10 & 0 & -10 \\ 0 & 0 & 10 & 8 \\ 10 & 0 & 0 & 8 \end{bmatrix}. \tag{8}$$

In that picture, the labels $i = 1,2,3$ correspond to the scaled negative unit vectors $-e_i/\varepsilon$, the labels $m + j = 4,5,6,7$ to the columns b_j of B. The nonpositive entries of A and B are allowed because a player's best-response payoff is always positive.

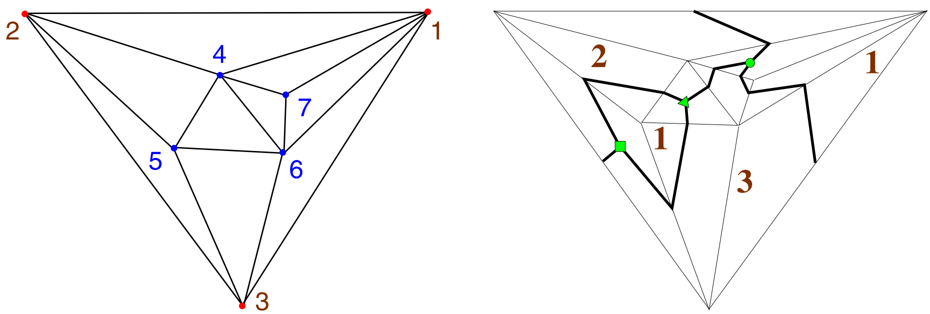

Fig. 1. Left: Schlegel diagram of $P_\varepsilon^\triangle$ for the example (8). Right: Subdivision into best-response regions of player 1, which completes the dual construction.

3 Subdividing Simplices into Best-Response Regions

The second step of our construction is the incorporation of player 1's best responses into the surface of $P_\varepsilon^\triangle$. Let $F(K,J)$ be a facet of $P_\varepsilon^\triangle$ as in (7). Consider the $m \times m$ matrix

$$[I_K\ A_J] = [e_{i_1} \cdots e_{i_k}\, a_{j_1} \cdots a_{j_{m-k}}] \qquad \text{if } K = \{i_1, \ldots, i_k\},\ J = \{j_1, \ldots, j_{m-k}\}, \tag{9}$$

that is, the columns of $[I_K\ A_J]$ are the columns e_i of the identity matrix I for $i \in K$ and the columns a_j of player 1's payoff matrix A for $j \in J$, where $A = [a_1 \cdots a_n]$. We write the unit simplex $\text{conv}\{e_1, \ldots, e_m\}$ in $\mathbb{R}^m$ as

$$\Delta(K,J) = \{z \in \mathbb{R}^K \times \mathbb{R}^J \mid z \geq \mathbf{0},\ \mathbf{1}^\top z = 1\}. \tag{10}$$

Proposition 4. *Let $(x,y) \in P \times Q - \{(\mathbf{0},\mathbf{0})\}$. Then (x,y) is a Nash equilibrium of (A,B) if and only if the vertex x of P corresponds to a facet $F(K,J)$ of $P_\varepsilon^\triangle$ so that $[I_K\ A_J]z = \mathbf{1}u$ for some $z = (z_K, z_J) \in \Delta(K,J)$ and some $u > 0$, and $y_J = z_J/u$, where y_J is y restricted to its support J.*

Proof. Because the game is nondegenerate, only vertices x of P can represent equilibrium strategies. Let $F(K,J)$ be the facet of $P_\varepsilon^\triangle$ that corresponds to x, where $K = \{i \mid x_i = 0\}$ and J is the set of best responses to x. Then y is a best response to x if and only if the support of y is J; suppose this holds, so that $Ay = A_J y_J$. In turn, x is a best response to y if and only if $(Ay)_i = 1$ whenever $i \notin K$, because $Ay \leq \mathbf{1}$. This is equivalent to $I_K s_K + A_J y_J = \mathbf{1}$ for a suitable slack vector $s_K \in \mathbb{R}^K$, $s_K \geq \mathbf{0}$. With $u = 1/(\sum_{i \in K} s_i + \sum_{j \in J} y_j)$ and $z = (s_K u, y_J u)$ this is equivalent to $z \in \Delta(K,J)$ and $[I_K\ A_J]z = \mathbf{1}u$ as claimed. □

Given a facet $F(K,J)$ of $P_\varepsilon^\triangle$ that corresponds to a vertex x of P, Prop. 4 states that x is part of a Nash equilibrium (x,y) if and only if there is a mixed strategy $z = (z_K, z_J) \in \Delta(K,J)$ so that all pure strategies of player 1 are best responses against z when the payoff matrix to player 1 is $[I_K\ A_J]$. Suitably scaled, z_K is a vector of slack variables, and z_J represents the nonzero part y_J of player 2's mixed strategy y. Nondegeneracy implies that z is completely mixed and hence in the interior of $\Delta(K,J)$.

The simplex $\Delta(K,J)$ has dimension $m-1$, like the face $F(K,J)$. The two simplices are in one-to-one correspondence via the canonical linear map

$$\alpha : \Delta(K,J) \to F(K,J), \qquad z \mapsto [M_K \ B_J]z, \tag{11}$$

where $M_K = (-1/\varepsilon \cdot I)_K$. This just says that α maps the vertices of $\Delta(K,J)$ (which are the unit vectors in $\mathbb{R}^m$) to the respective vertices of $F(K,J)$, and preserves convex combinations.

We subdivide $\Delta(K,J)$ into polyhedral *best response regions* $\Delta(K,J)(i)$ for the strategies $i = 1,\ldots,m$ of player 1, using the payoff matrix $[I_K \ A_J]$. That is (see [13] or [15]), $\Delta(K,J)(i)$ is the set of mixed strategies z so that i is a best response to z, so for $1 \le i \le m$,

$$\Delta(K,J)(i) = \{z \in \Delta(K,J) \mid ([I_K \ A_J]z)_i \ge ([I_K \ A_J]z)_k \text{ for all } k = 1,\ldots,m\}. \tag{12}$$

We say z in $\Delta(K,J)$ has *label* i if $z \in \Delta(K,J)(i)$, and correspondingly $\alpha(z)$ in $F(K,J)$ has label i if z has label i.

This *dual construction* [14] labels every point on the surface of $P_\varepsilon^\triangle$. The labeling is unique because the payoffs to player 1, and the map α, only depend on the vertices of the respective facets, so the labels agree for points that belong to more than one facet. For the game in (8), this labeling is shown in the right picture of Fig. 1.

As a consequence of Prop. 4, we obtain that the equilibria of the game correspond to the points on the surface of $P_\varepsilon^\triangle$ that have all labels $1,\ldots,m$. We call such points *completely labeled*. The three equilibria of the game (8) are marked by a small square, triangle and circle in Fig. 1. Not all facets of $P_\varepsilon^\triangle$ contain such a completely labeled point, if the corresponding vertex x of P is not part of a Nash equilibrium.

"Completely labeled" now means "all strategies of player 1 appear as labels". What happened to the strategies of player 2? They correspond to the vertices of $P_\varepsilon^\triangle$. They are automatically best responses when considering the facets of $P_\varepsilon^\triangle$, and they are the only strategies that player 2 is allowed to use, apart from the slacks $i \in K$, when subdividing a facet $F(K,J)$ into the labeled regions $\alpha(\Delta(K,J)(i))$ for the labels $i = 1,\ldots,m$.

4 Visualizing the Index and the LH Algorithm

The described dual construction, the labeled subdivision of the surface of $P_\varepsilon^\triangle$, visualizes all equilibria of an $m \times n$ game in a geometric object of dimension $m-1$. Figure 2 also shows the index of an equilibrium as the orientation in which the labels $1,\ldots,m$ appear around the point representing the equilibrium, here counterclockwise for index $+1$, and clockwise for index -1. The artificial equilibrium is the completely labeled point $M(\mathbf{1}/m)$ (see (11) with $J = \emptyset$) of the facet F_0 of $P_\varepsilon^\triangle$, which has negative index. This facet should be seen as the underside of the "flattened" view of $P_\varepsilon^\triangle$ given by the Schlegel diagram, so the dashed border of F_0 in Fig. 2 is to be identified with the border of the large triangle.

Our goal is to formalize the orientation of a completely labeled point in the dual construction, and to show that it agrees with the index in Def. 1. A nonsingular $m \times m$ matrix C has *positive orientation* if $\det(C) > 0$, otherwise *negative orientation*.

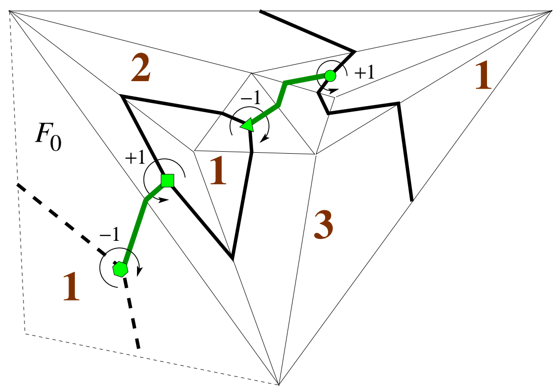

Fig. 2. Indices of equilibria as positive or negative orientations of the labels, and LH paths for missing label 1. The facet on the left with dashed border indicates the flapped-out "underside" of the Schlegel diagram, the facet F_0 of $P_\varepsilon^\triangle$.

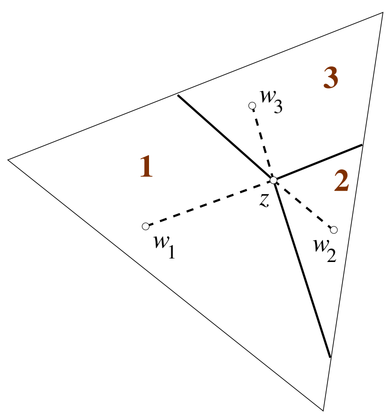

Fig. 3. Points w_1, w_2, w_3 in $\Delta(K,J)$, here for $K = \emptyset$, so that (13) holds

Let $\alpha(z)$ be a completely labeled point of a facet $F(K,J)$ of $P_\varepsilon^\triangle$. We first consider points $w_1, \ldots, w_m$ so that w_i belongs only to the best-response region $\Delta(K,J)(i)$ for $1 \le i \le m$. More specifically, we want that for suitable $s_i \ge 0$, $t_i > 0$,

$$[I_K\ A_J]w_i = \mathbf{1}s_i + e_i t_i, \tag{13}$$

that is, player 1's payoff against w_i is $s_i + t_i$ for his pure strategy i, and a smaller constant s_i for all other pure strategies $k \ne i$. Such points w_i exist, by extending the line through the completely labeled point z defined by the $m-1$ labels $k \ne i$ into the region $\Delta(K,J)(i)$, as shown in Fig. 3. For $i \in K$, we can simply choose $w_i = e_i$ to obtain (13), a case that is not shown in Fig. 3.

Let $W = [w_1 \cdots w_m]$. We want to show that W has the same orientation as $[I_K\ A_J]$. Because of (13), $[I_K\ A_J]W = C + [\mathbf{1}s_1 \cdots \mathbf{1}s_m]$ for the diagonal matrix C with entries

$c_{ii} = t_i > 0$ and $c_{ij} = 0$ for $i \neq j$. By (3), $C + [\mathbf{1}s_1 \cdots \mathbf{1}s_m]$ has positive determinant, so that $\det[I_K\ A_J]$ and $\det(W)$ have the same sign, as claimed.

We take the orientation of the matrix $D = [\alpha(w_1) \cdots \alpha(w_m)]$ as the orientation of the equilibrium represented by $\alpha(z)$. By (11), that matrix is $D = [M_K\ B_J]\,W$. Its orientation $\operatorname{sign}(\det(D))$ is the sign of $\det[M_K\ B_J]\det(W)$, so that

$$\operatorname{sign}(\det(D)) = \operatorname{sign}(\det[M_K\ B_J])\operatorname{sign}(\det[I_K\ A_J]). \tag{14}$$

Let $L = \{1,\ldots,m\} - K$, which is the support of the vertex x of P that corresponds to the facet $F(K,J)$. We can assume that $K = \{1,\ldots,k\}$, because any transposition of player 1's strategies alters the signs of both determinants on the right-hand side of (14). Then

$$\begin{aligned}\operatorname{sign}(\det(D)) &= \operatorname{sign}((-1/\varepsilon)^k \det(B_{LJ}))\ \operatorname{sign}(\det(A_{LJ}))\\ &= (-1)^{m-|L|}\operatorname{sign}(\det(B_{LJ}))\ \operatorname{sign}(\det(A_{LJ}))\\ &= (-1)^{m-1}(-1)^{|L|+1}\operatorname{sign}(\det(B_{LJ}))\ \operatorname{sign}(\det(A_{LJ})).\end{aligned}$$

Consequently, $\operatorname{sign}(\det(D))$ agrees with the index of the equilibrium when m is odd, and is the negative of the index when m is even. The artificial equilibrium corresponds to the center point of F_0, which has orientation $(-1)^m$. The orientation of the artificial equilibrium should always be -1, so it has to be multiplied with -1 when m is even. Hence, relative to the orientation of the artificial equilibrium, $\operatorname{sign}(\det(D))$ is exactly the index of the equilibrium under consideration, as claimed.

We mention very briefly an interpretation of the LH algorithm with the dual construction, as illustrated in Fig. 2; for details see [14]. This can only be done for missing labels of player 1, because player 2 is always in equilibrium. For missing labels of player 2 one has to exchange the roles of the two players (the dual construction works either way). The original LH path starts from $(\mathbf{0},\mathbf{0})$ in $P \times Q$ by dropping a label, say label 1, in P. This corresponds to dropping label 1 from the artificial equilibrium given by the center of F_0. It also reaches a new vertex of P, which in $P_\varepsilon^\triangle$ means a change of facet. This means a change of the normal vector of that facet, which is an invisible step in the dual construction because the path is at that point on the joint face of the two facets. Traversal of an edge of Q is represented by traversing the line segment in a face $P_\varepsilon^\triangle$ that has all $m-1$ labels except for the missing label. That line segment either ends at an equilibrium, or else reaches a new facet of $P_\varepsilon^\triangle$. The path then (invisibly) changes to that facet, followed by another line segment, and so on. Algorithmically, the LH pivoting steps are just like for the path on $P \times Q$, so nothing changes.

Figure 2 also illustrates why the endpoints of LH paths have opposite index: Along the path, the $m-1$ labels that are present preserve their orientation around the path, whereas the missing label is in a different direction at the beginning and at the end of the path. In Fig. 2, an LH path from a -1 to a $+1$ equilibrium with missing label 1 has label 2 on the left and label 3 on the right. This intuition of Shapley's result Prop. 2(e) [13] can be given without the dual construction (see [12, Fig. 1]), but here it is provided with a figure in low dimension.

5 Proof Sketch of Theorem 3

In this section, we give an outline of the proof of Theorem 3 with the help of the dual construction. We confine ourselves to an equilibrium (x,y) of index $+1$ that uses all m strategies of player 1, which captures the core of the argument. Hence, the facet $F(K,J)$ of $P_\varepsilon^\triangle$ that corresponds to the fully mixed strategy x of player 1 has $K = \emptyset$. The m best responses of player 2 are $j \in J$, which define the payoff vectors b_j as points in $\mathbb{R}^m$, and an $m \times m$ submatrix B_J of B.

For player 1 and player 2, we will construct three $m \times m$ matrices A', A'', A''' and B', B'', B''', respectively, so that the extended game G' in Theorem 3 is defined by the two $m \times (n + 3m)$ payoff matrices $[A\,A'\,A''\,A''']$ and $[B\,B'\,B''\,B''']$.

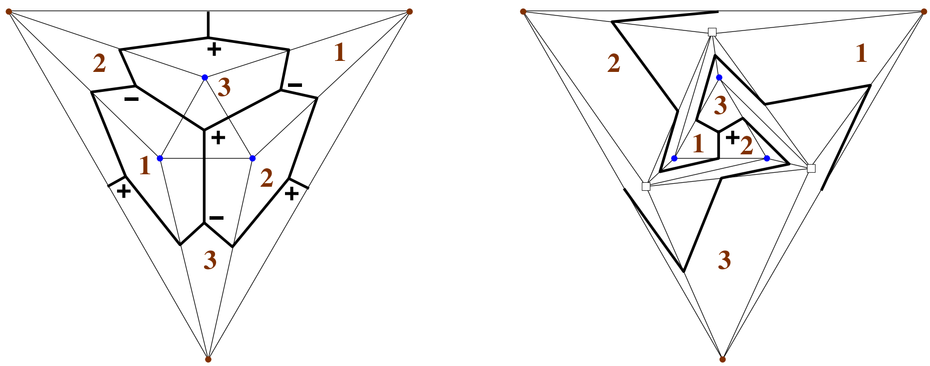

Fig. 4. Left: Dual construction for a 3×3 coordination game, which has four equilibria with index $+1$ and three with index -1. Right: Adding three strategies for the column player (small white squares) so that only the completely mixed equilibrium remains, see (15).

An example is the 3×3 coordination game where both players' payoffs are given by the identity matrix. The game has seven equilibria: three pure-strategy equilibria and the completely mixed equilibrium, all of index $+1$, and three equilibria where each player mixes two strategies, with index -1. The left picture in Fig. 4 shows the dual construction for this game. The index of each equilibrium point is indicated by its sign, given by the orientation of the labels $1,2,3$ around the point. The completely mixed equilibrium is on the central triangle with facet F whose vertices are the three columns of B. We want to make this equilibrium unique by adding strategies. In this case, we need only the matrices A'' and B'', for example as follows:

$$[A\,A''] = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 1 & 0 \end{bmatrix}, \quad [B\,B''] = \begin{bmatrix} 11 & 10 & 10 & 12 & 8.9 & 10 \\ 10 & 11 & 10 & 10 & 12 & 8.9 \\ 10 & 10 & 11 & 8.9 & 10 & 12 \end{bmatrix}. \qquad (15)$$

The dual construction for the game in (15) is shown on the right in Fig. 4. As desired, only the original facet F has an equilibrium point, which is now unique. It is also clear that its index must be $+1$, because otherwise it would not be possible to "twist" the best response regions $1,2,3$ outwards to meet the labels at the outer vertices.

In this example, the columns of B'' span a simplex (with vertices indicated by small white squares in Fig. 4), whose projection to F_0 in the Schlegel diagram contains the original facet F as a subset. In fact, the columns of B'' are first constructed as points in the hyperplane defined by F so that they define a larger simplex than F. Subsequently, these new points are moved slightly to the origin, so that F re-appears in the convex hull: Note that in (15), the normal vector for the hyperplane through the columns of B'' is $\mathbf{1}$, but its scalar product with these columns is 30.9 and not 31 like for the columns of B (the matrix B is the identity matrix with 10 added to every entry).

In the general construction, several complications have to be taken care of. First, the original game may have additional columns that are not played in the considered equilibrium. The example makes it clear that this is a minor complication: Namely, the simplex spanned by the columns of B'' can be magnified, while staying in the hyperplane just below F, so that the convex hull of these columns and of the negative unit vectors contains all unused columns of B in its interior.

A second complication is that the labels $1,\ldots,m$ of the best-reponse regions given by A may not correspond to the vertices of F as they do in Fig. 4. For example, two of the vertices of the triangle in Fig. 3 have label 1, one vertex has label 3, and no vertex has label 2. The first matrix B' in the general construction is designed so that each label $1,\ldots,m$ appears at exactly one vertex. Namely, consider the simplex spanned by the points $w_1,\ldots,w_m$ that are used to represent the best-response regions $1,\ldots,m$ around the equilibrium point z (after these points have been mapped into F via α). If this simplex is "blown up" around z while staying in the same hyperplane, it eventually contains the original unit simplex. Let $v_1,\ldots,v_m$ be the vertices of this blown-up simplex, as shown in Fig. 5. After the mapping α, the corresponding points will be

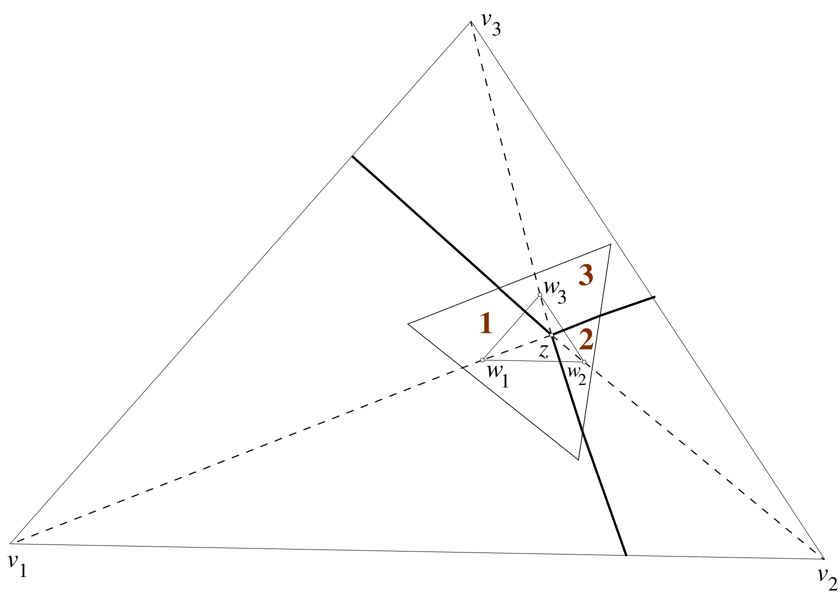

Fig. 5. Points $v_1,\ldots,v_m$ as expanded points $w_1,\ldots,w_m$ along the dashed lines around z in the same hyperplane, so that $\mathrm{conv}\{v_1,\ldots,v_m\}$ contains the original simplex $\Delta(K,J)$

in the hyperplane defined by F and define a simplex that contains F as a subset. We merely move these points $\alpha(v_1),\ldots,\alpha(v_m)$ slightly towards to the origin, which defines the matrix B'. The corresponding payoffs A' to player 1 are given by the diagonal matrix $[e_1 t_1 \cdots e_m t_m]$ with the payoffs $t_1,\ldots,t_m$ given as in (13). We could add an arbitrary constant $\mathbf{1}s_i$ to the ith column of A' (for each i) without changing the subdivision into best-response regions of the simplex defined by B'. Hence, if B' was still in the same hyperplane as F, that subdivision would coincide with the subdivision of F, which it still does after moving B' slightly inwards. From then on, we consider the simplex by spanned B' instead of F, which then looks essentially like in the special case of Fig. 15 because the corresponding matrix A' is a diagonal matrix.

We also use two increasingly larger simplices defined by B'' and B''', with identity matrices A'' and A'''. In the resulting construction, each pair of matrices (M,B'''), (B''',B''), (B'',B') and (B',B_J) (where the columns of M and B_J are the vertices of F_0 and F, respectively) defines a pair of simplices whose convex hull is a "distorted prism". These distorted prisms are stacked on top of each other, with the points of intermediate layers spread outwards in parallel hyperplanes to maintain a convex set. In the projected Schlegel diagram, each simplex is contained in the next.

The missing matrices B'' and B''' are constructed using the following theorem of [8]: *Every matrix with positive determinant is the product of three P-matrices.* A P-matrix is a matrix P such that every principal minor P_{SS} of P (where S is an arbitray set of rows and the same set of columns) has positive determinant. The P-matrices are useful for "stacking distorted prisms" because of the following property:

Proposition 5. *Let $P = [p_1 \cdots p_m]$ be a* P*-matrix where each column p_i is scaled such that $\mathbf{1}^\top p_i = 2$. Let X be the convex hull of the vectors p_i and the unit vectors e_i for $1 \le i \le m$. Assume X is a simplicial polytope, if necessary by slightly perturbing P. Let a facet of X have label i if it has p_i or e_i has a vertex. Then the only facets of X that have all labels* $1,\ldots,m$ *are those spanned by $p_1,\ldots,p_m$ and by $e_1,\ldots,e_m$.*

Proposition 5 may be a novel observation about P-matrices. It can be proved using a parity argument: Additional completely labeled facets would have to come in pairs of opposite orientation, and a negatively oriented facet contradicts the P-matrix property.

Consequently, a distorted prism X defined by the columns of the identity matrix I and a P-matrix P (scaled as in Prop. 5) has no completely labeled facet other than its two "end" facets defined by I and P. If Q is another such P-matrix, the same observation holds for I and Q, and consequently for P and PQ because it does not change under affine transformations. Finally, for another P-matrix R, we see that PQ and PQR define prisms that have no completely labeled "side" facets, either. According to said theorem of [8], PQR can represent an arbitrary matrix with positive determinant.

The stack of prisms that we want to generate should start at the convex hull M of the negative unit vectors used in (11). We move these vectors in the direction $\mathbf{1}$ until it crosses the origin, so that the resulting matrix, which we call N, has opposite orientation to M. As shown in Section 4, N has therefore the same orientation as the matrix D in (14) and hence as B'. Therefore, $N^{-1}B'$ has positive determinant and can be represented as a product PQR of three P-matrices, so that $NPQR = B'$. Then the additional matrices are given by $B'' = NPQ$ and $B''' = NP$.

We have to omit details for reasons of space. We conclude with the following intuition why we use two matrices B'' and B''' rather than just a single one. In Fig. 4, the columns of B'', which is the only matrix used, are indicated by the white squares, but these show only the projection in the Schlegel diagram. Their (invisible) distances from the origin are very important because they determine the facets spanned by the columns of B'' and of B. Essentially, B'' represents a single intermediate step when "twisting" F by 180 degrees towards the outer triangle, and A'' is a matrix of suitably ordered unit vectors. It is not clear if this can be done in higher dimensions. With two intermediate sets of points B'' and B''', their exact relative position is not very relevant when using P-matrices, due to Prop. 5, and one can use identity matrices A'' and A'''.

References

1. Bárány, I., Vempala, S., Vetta, A.: Nash equilibria in random games. In: Proc. 46th FOCS, pp. 123–131 (2005)
2. Chen, X., Deng, X.: Settling the complexity of 2-player Nash-equilibrium. In: Proc. 47th FOCS, pp. 261–272 (2006)
3. Conitzer, V., Sandholm, T.: Complexity results about Nash equilibria. In: Proc. 18th International Joint Conference on Artificial Intelligence (IJCAI), pp. 765–771 (2003)
4. Demichelis, S., Germano, F.: On the indices of zeros of Nash fields. Journal of Economic Theory 94, 192–217 (2000)
5. Gilboa, I., Zemel, E.: Nash and correlated equilibria: some complexity considerations. Games and Economic Behavior 1, 80–93 (1989)
6. Govindan, S., Wilson, R.: Uniqueness of the index for Nash equilibria of two-player games. Economic Theory 10, 541–549 (1997)
7. Hofbauer, J.: Some thoughts on sustainable/learnable equilibria. In: 15th Italian Meeting on Game Theory and Applications, Urbino, Italy, July 9–12 (2003) (accessed on October 29, 2007), `http://www.econ.uniurb.it/imgta/PlenaryLecture/Hofbauer.pdf`
8. Johnson, C.R., Olesky, D.D., van den Driessche, P.: Matrix classes that generate all matrices with positive determinant. SIAM J. Matrix Anal. Appl. 25, 285–294 (2003)
9. Lemke, C.E., Howson Jr., J.T.: Equilibrium points of bimatrix games. Journal of the Society for Industrial and Applied Mathematics 12, 413–423 (1964)
10. Myerson, R.B.: Sustainable equilibria in culturally familiar games. In: Albers, W., et al. (eds.) Understanding Strategic Interaction: Essays in Honor of Reinhard Selten, pp. 111–121. Springer, Heidelberg (1997)
11. Papadimitriou, C.H.: On the complexity of the parity argument and other inefficient proofs of existence. Journal of Computer and System Sciences 48, 498–532 (1994)
12. Savani, R., von Stengel, B.: Exponentially many steps for finding a Nash equilibrium in a bimatrix game. In: Proc. 45th FOCS, pp. 258–267 (2004)
13. Shapley, L.S.: A note on the Lemke–Howson algorithm. Mathematical Programming Study 1: Pivoting and Extensions, pp. 175–189 (1974)
14. von Schemde, A.: Index and Stability in Bimatrix Games. Lecture Notes in Economics and Mathematical Systems, vol. 1853. Springer, Berlin (2005)
15. von Stengel, B.: Computing equilibria for two-person games. In: Aumann, R.J., Hart, S. (eds.) Handbook of Game Theory, vol. 3, pp. 1723–1759. North-Holland, Amsterdam (2002)
16. Ziegler, G.M.: Lectures on Polytopes. Springer, New York (1995)