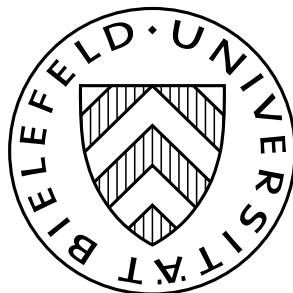


July 2006

A Simplicial Algorithm Approach to Nash Equilibria in Concave Games

Claus-Jochen Haake and Francis Edward Su



IMW · Bielefeld University
Postfach 100131
33501 Bielefeld · Germany



email: imw@wiwi.uni-bielefeld.de
<http://www.wiwi.uni-bielefeld.de/~imw/Papers/showpaper.php?382>
ISSN: 0931-6558

A Simplicial Algorithm Approach to Nash Equilibria in Concave Games*

Claus-Jochen Haake[†] Francis Edward Su[‡]

July 2006

Abstract

In this paper we demonstrate a new method for computing approximate Nash equilibria in n -person games. Strategy spaces are assumed to be represented by simplices, while payoff functions are assumed to be concave. Our procedure relies on a simplicial algorithm that traces paths through the set of strategy profiles using a new variant of Sperner's Lemma for labelled triangulations of simplotopes, which we prove in this paper. Our algorithm uses a labelling derived from the satisficing function of Geanakoplos (2003) and can be used to compute approximate Nash equilibria for payoff functions that are not necessarily linear. Finally, in bimatrix games, we can compare our simplicial algorithm to the combinatorial algorithm proposed by Lemke & Howson (1964).

Key words: Simplicial algorithm, Nash equilibria, strategy labelling

MSC 2000 classification codes: Primary: 91B50 Secondary: 91A10

OR / MS classification: Primary: Noncooperative, Combinatorics Secondary: Fixed points

*The authors gratefully acknowledge partial support by NSF Grant DMS-0301129 (Su).

[†]Institute of Mathematical Economics, Bielefeld University, P.O. Box 100131, 33501 Bielefeld, Germany; e-mail: CHaake@wiwi.uni-bielefeld.de.

[‡]Department of Mathematics, Harvey Mudd College, Claremont, CA 91711, USA; e-mail: su@math.hmc.edu.

1 Introduction

There is a well-established equivalence between the Brouwer and Kakutani fixed point theorems in the sense that proofs relying on one of the theorems can be proved using either version (see, e.g., Border (1985)). In particular, existence of Nash equilibria for the class of non-cooperative n -person games with compact, convex strategy sets and continuous, quasi-concave payoffs can be proved by using the Brouwer fixed point theorem. In a recent paper, Geanakoplos (2003) proves the existence of Nash equilibria using the Brouwer theorem by introducing a “satisficing improvement *function*”, which replaces the usual best-reply correspondence normally associated to the Kakutani fixed point theorem. For each agent, this satisficing function takes a strategy profile to a strategy of this agent that is a better response (though not necessarily best). Nonetheless, it turns out that fixed points of this function are Nash equilibria of the underlying game.

In a different strand of literature, so-called *simplicial algorithms* are used to compute approximate fixed points of continuous functions in the following way. First, the underlying domain is triangulated, i.e., it is decomposed into simplices in some nice way. Next, each vertex (or each face) in this triangulation is assigned a *label* that reflects the structure of the given function. Finally, by using a path-following argument, one detects a sequence of cells of the decomposition that leads to a specific final cell, which has the property that points in that cell approximately satisfy the fixed point condition. In this sense, a simplicial algorithm finds an approximate fixed point, and such algorithms adapted by *homotopy* methods can converge to an actual fixed point. See Yang (1999) for a recent survey of this literature.

Similar mechanisms have been frequently used to compute approximate Nash equilibria. For example, van der Laan & Talman (1982) introduce a simplicial algorithm that is used to compute Nash equilibria of games with multilinear payoff functions. Herings & van den Elzen (2002) provide an algorithm that mimics Harsanyi’s (non-constructive) tracing procedure and finds ε -Nash equilibria.

In this paper we demonstrate a new method for computing (approximate) Nash equilibria for a large class of n -person games by combining Geanakoplos’s (2003) approach with simplicial algorithm techniques. We use a new

variant of Sperner’s lemma that we prove in this paper. A constructive proof of our *Simplotopal Sperner Lemma* (Theorem 2.1) provides a basis for a simplicial algorithm that traces paths through the set of strategy profiles (a geometric object known as a *simplotope*) to locate a Nash equilibrium.

While the computation of Nash equilibria can be accomplished using the usual Sperner lemma (e.g., see Todd (1976)) or other combinatorial theorems (e.g., van der Laan & Talman (1982), Freund (1986)), the labelling function in our simplotopal Sperner lemma differs from these other approaches in that it allows us to disentangle the preferences of the various players into various coordinates in the labelling. Moreover, it naturally meshes with the mapping of Geanakoplos (2003) to produce an algorithm that works for payoff functions which are not necessarily linear. This contrasts with other algorithms for computing Nash equilibria that only work in the linear case, e.g., Lemke & Howson (1964) or van der Laan & Talman (1982). Through the choice of a *flag of faces* of the simplotope, our algorithm also allows for the possibility of locating many different Nash equilibria.

2 A Simplotopal Sperner Lemma

A *simplotope* S is a product of simplices (of possibly varying dimension). We write $S = \prod_{i=1}^n S^i$, where S^i is the i -th factor. For our purposes later S^i will represent the strategy set of player i : if player i has d^i “pure” strategies e^1, \dots, e^{d^i} in \mathbb{R}^{d^i} , then S^i is their convex hull and therefore a $(d^i - 1)$ -simplex. Then $S = \prod_{i=1}^n S^i$ is the set of strategy profiles for all the players. Every point $s = (s^1, \dots, s^n)$ of S is a choice of a (mixed) strategy $s^i \in S^i$ for every player; moreover, each s^i is a vector $(s_1^i, \dots, s_{d^i}^i)$ that represents a convex combination of player i ’s pure strategies, hence $\sum_j s_j^i = 1$ and $s_j^i \geq 0$. Define the *carrier* of a vector s^i in S^i to be the set $C(s^i) = \{j : s_j^i > 0\}$. In other words, the carrier of a mixed strategy is the set of all the pure strategies that it involves.

We shall be interested in certain labelled triangulations of the simplotope S . A *triangulation* T of S is a collection of simplices satisfying: (i) if σ is a simplex in T , then any face of σ is in T , (ii) the union of the simplices in T is S , and (iii) the intersection of any two simplices σ, τ in T is a face of both σ and τ (possibly the empty face). Loosely speaking, a triangulation of S is

a subdivision of S into simplices that meet in a nice way, i.e., face-to-face.

Let $L^i = \{1, \dots, d^i\}$ be a set of integers, one for each of player i 's pure strategies. We shall use the vectors in the product $L = \prod_{i=1}^d L^i$ as a set of labels. Note that each label corresponds to a vector of pure strategies, one for each player. If T is a triangulation of S , a *strategy-labelling* of T is a function ℓ taking each vertex of T to an element of the label set $L = \prod_{i=1}^n L^i$.

Consider the projections $\pi^i : S \rightarrow S^i$ defined by $\pi^i(s^1, \dots, s^n) = s^i$. There is a similar projection $\pi^i : L \rightarrow L^i$ defined by $\pi^i(k_1, \dots, k_n) = k_i$ (for convenience we give it the same name even though it is a different projection). Let $\ell^i : S \rightarrow L^i$ be the i -th coordinate function of ℓ , i.e., $\ell(s) = (\ell^1(s), \ell^2(s), \dots, \ell^n(s))$. Hence $\ell^i(s) = \pi^i(\ell(s))$.

We shall call a strategy-labelling ℓ a *simplotopal Sperner labelling* if $s_{\ell^i(s)}^i > 0$ holds for each i . In other words, for each i , $\ell^i(s)$ must be chosen from one of the labels in the carrier of $\pi^i(s)$. Put still another way, ℓ is a simplotopal Sperner labelling if when all coordinates of $s = (s^1, \dots, s^n)$ are fixed except for s^i (making ℓ a function on S^i), then the i -th coordinate labels $\ell^i(s)$ form a Sperner labelling over S^i .

A *full cell* in a triangulation T with labelling ℓ is a simplex σ in T all of whose vertices carry distinct label vectors. We say that σ is *full* in (T, ℓ) .

Given σ in T , let $L(\sigma)$ be the label set of the vertices of σ , and let $L_i(\sigma)$ be the i -th coordinates of the label vectors of σ . Hence $L_i(\sigma) \subseteq L^i$.

If a full cell satisfies $L_i(\sigma) = L^i$ for *every* coordinate i , then we say that σ is *coordinate-full* or *c-full* in (T, ℓ) . We prove the following theorem.

Theorem 2.1 (A Simplotopal Sperner Lemma). *Suppose that a simplotope S with triangulation T has a simplotopal Sperner-labelling ℓ . Then there exists a c-full cell σ in (T, ℓ) .*

We immediately note the difference between our labelling function ℓ and the labelling functions of Freund (1986) and van der Laan & Talman (1982), who labelled vertices in the simplotope triangulation by *facets* of S , rather than by the main vertices of S . We remark later on the consequences of this for the resulting algorithm to find Nash equilibria.

Example 2.2. If player 1 has $d^1 = 2$ strategies, and player 2 has $d^2 = 2$ strategies, then a point in S^i is a 2-tuple (s_1^i, s_2^i) such that $s_1^i + s_2^i = 1$ and $s_1^i, s_2^i \geq 0$. Thus each S^i may be viewed as a line segment with endpoints $(1, 0)$ and right endpoint $(0, 1)$. Then the set of strategy profiles S is the product of two line segments, i.e., a square. A generic point in S is of the form $((s_1^1, s_2^1), (s_1^2, s_2^2))$.

For a labelling, let $L^1 = \{1, 2\}$, $L^2 = \{1, 2\}$. Then L contains four labels that are ordered pairs of strategies: $(1, 1), (1, 2), (2, 1), (2, 2)$. For simplicity, we abbreviate these labels: $L = \{11, 12, 21, 22\}$.

A simplotopal Sperner labelling of a triangulation of S is a labelling of each point v by vectors whose vector coordinates specify a factor S^i in which $v = (v_1, \dots, v_n)$ has a non-zero coordinate v_i . See Figure 1. Hence all points along the left side (respectively, right side) of S must have first coordinate 1 (respectively, 2). Similarly, all points along the bottom (respectively, top) side of the square must have second coordinate 1 (respectively, 2).

As we claim in Theorem 2.1, a simplotopal Sperner labelling always contains a c-full cell σ , which means that the cell has all its vertex labels distinct, and for each player i , the i -th coordinates of all the labels of σ must represent the full set of pure strategies of player i . We indeed see such a c-full cell, whose labels are $\{12, 21, 22\}$; notice that all the labels are distinct, and labels 1 and 2 appear among the first coordinates of the labels while labels 1 and 2 appear among the second coordinates of the labels. \square

Example 2.3. As another example, if player 1 has $d^1 = 3$ strategies, and player 2 has $d^2 = 2$ strategies, then the set of strategy profiles S is the product of a 2-simplex and a 1-simplex, i.e., a prism. Here $L^1 = \{1, 2, 3\}$, $L^2 = \{1, 2\}$, and L now contains six labels (in abbreviated form): $L = \{11, 21, 31, 12, 22, 32\}$. See Figure 1. Let T be a triangulation of S . A simplotopal Sperner labelling of T would label each vertex (v_1, v_2) of T by a vector (l^1, l^2) (where $l^i \in L^i$ for each i) in such a way that l^i is in the carrier $C(v_i)$ for $i = 1, 2$. In Figure 1, this means that vertices on the bottom triangular face will have $l^2 = 1$, on the top triangular face will have $l^2 = 2$, and on each vertical edge l^1 is constant and equal to 1, 2, or 3. Then $\ell(e^i, e^j) = (i, j)$ for each $i \in L^1$ and $j \in L^2$. Furthermore, we shall show below that any vertex will be labelled by the vector label of one of the

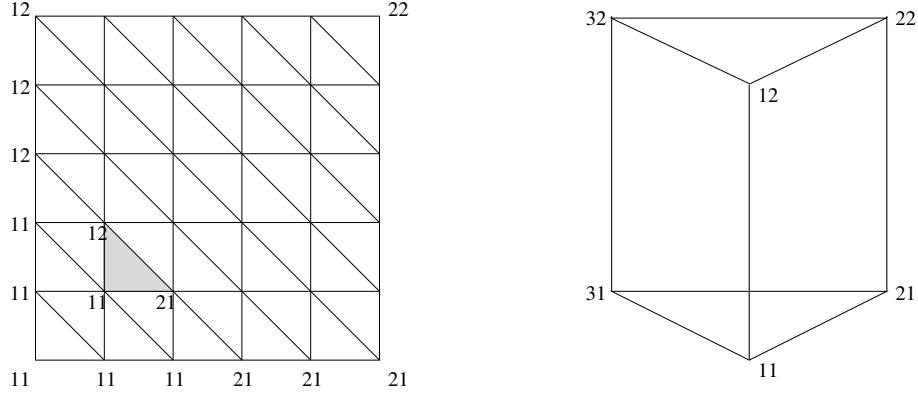


Figure 1: Left: A simplotopal Sperner-labelled triangulation of a square. All vertices without labels are labelled 22. A full cell is shaded. Right: Vertex labelling of a prism.

(e^i, e^j) that span the face that carries that vertex, e.g., vertices on the face $\{32, 12, 31, 11\}$ must be labelled by one of the labels $\{32, 12, 31, 11\}$. \square

First we prove a lemma. A *polytopal Sperner labelling* (De Loera, Peterson & Su 2002) of a polytope P is a labelling ℓ such that (i) all vertices of P have distinct labels, and (ii) if a vertex v lies in the interior of a face of P , then $\ell(v) = \ell(w)$ for some vertex w of S that spans that face.

Lemma 2.4. *A simplotopal Sperner-labelling ℓ of a triangulation of simpletope S is a polytopal Sperner labelling.*

Proof. First we note that the simplotopal Sperner labelling ℓ labels all vertices of S differently. To show this, suppose that $w = (w^1, \dots, w^n) = (e^{k_1}, \dots, e^{k_n})$, where e^{k_i} is a unit vector in \mathbb{R}^{d_i} with 1 in the k_i -th coordinate. Then for each i , there is exactly one element in the carrier $C(w^i)$ and it must be k_i . Thus the label of w is completely determined by the simplotopal Sperner lemma condition; it must be $\ell(w) = (k_1, \dots, k_n)$.

Next we show that if a vertex v lies in a face of S , then $\ell(v) = \ell(w)$ for some vertex w of S that spans that face. The vertex $v = (v^1, \dots, v^n)$ lies on some minimal face F of S , which is the product of n faces in the factors S^i of S . In fact, $F = F_1 \times \dots \times F_n$ where each F_i is the convex hull of $\{e^k : k \in C(v^i)\}$. Moreover, we may write v^i as the unique strict convex

combination: $v^i = \sum_{k \in C(v^i)} v_k^i e^k$. From the simplotopal Sperner labelling we know that $\ell^i(v)$ is in $C(v^i)$ for each i , say it is k_i . Then $\ell(v) = \ell(w)$ where $w = (e^{k_1}, \dots, e^{k_n})$ is a vertex of S and w is clearly a vertex of the face F . \square

To prove Theorem 2.1, we can now use the Polytopal Sperner Lemma of De Loera, Peterson & Su (2002):

Theorem 2.5 (Polytopal Sperner Lemma). *Any triangulation of a d -dimensional polytope P with n vertices with a polytopal Sperner labelling must contain at least $n - d$ non-degenerate full cells.*

Proof of Theorem 2.1. Let ℓ be a simplotopal Sperner labelling of a triangulation of S . By the Polytopal Sperner Lemma of De Loera, Peterson & Su (2002), there must be a full cell σ in (T, ℓ) . Moreover, it is *non-degenerate* in the sense that under the natural simplicial map $f : T \rightarrow T$ that takes a vertex v of the triangulation to the vertex of S labelled $\ell(v)$, $f(\sigma)$ has volume greater than zero.

On the other hand, the only simplices of S that span volume are ones for which $L_i(\sigma) = L^i$, otherwise $f(\sigma)$ would lie in some proper face of S . \square

We remark that one of the proofs of De Loera, Peterson & Su (2002) is a constructive path-following argument. The constructive aspect of this proof will be important for us later, so we give a sketch of how this path-following idea will find a full cell. While the discussion in the next few paragraphs may seem technical, the main point to remember is that the labelling provides a well-defined way to trace a path through the simplices of the triangulation, in a manner that can be readily implemented computationally.

The proof adapted to our context relies on choosing a *flag of faces* of the given simplotope S , which is an increasing sequence of faces $F_0 \subset F_1 \subset \dots \subset F_d = S$ in which each F_i is a face of dimension i . Call this flag \mathcal{F} . For any simplex σ , let $L(\sigma)$ denote the set of distinct labels of vertices of σ , and let $L(F_i)$ similarly denote the label set of vertices of face F_i .

We define the *flag graph* $G_{\mathcal{F}}$ as follows. A k -dimensional simplex σ is a node of the flag graph, if and only if it is of one of the following types:

1. σ is carried by the k -face F_k and

$$|L(\sigma) \cap L(F_j)| = j + 1 \quad \text{for all } 1 \leq j \leq k.$$

2. σ is carried by the $(k+1)$ -face F_{k+1} and

$$|L(\sigma) \cap L(F_j)| = j + 1 \quad \text{for all } 1 \leq j \leq k.$$

3. σ is carried by the k -face F_k and

$$\begin{aligned} |L(\sigma) \cap L(F_k)| &= k & \text{and} \\ |L(\sigma) \cap L(F_j)| &= j + 1 & \text{for all } 1 \leq j \leq k - 1. \end{aligned}$$

4. σ is carried by the k -face F_k and there is a J such that

$$\begin{aligned} |L(\sigma) \cap L(F_k)| &= k + 1, \\ |L(\sigma) \cap L(F_j)| &= j + 2 & \text{for all } J \leq j \leq k - 1 & \quad \text{and} \\ |L(\sigma) \cap L(F_j)| &= j + 1 & \text{for all } 1 \leq j < J \end{aligned}$$

Two nodes σ and τ carried by F_k are adjacent in $G_{\mathcal{F}}$, if (as simplices) σ is a facet of τ . Nodes σ carried by F_{k-1} and τ carried by F_k are adjacent, if σ is a facet of τ and σ is of type 1. There are no other adjacencies.

De Loera, Peterson & Su (2002) show that under these adjacency rules, all the nodes of the flag graph $G_{\mathcal{F}}$ have degree 1 or 2, and it has degree 1 if and only if it is a 1-simplex of $G_{\mathcal{F}}$ or a cell of Type 1 above. Note that cells of Type 1 are precisely the c-full cells, and there are an odd number of full 1-simplices in F_1 . Thus the flag graph consists of paths that must match up at least one of the full 1-simplices in F_1 with a c-full cell in F_d . These paths form the basis of the associated constructive simplicial algorithm. Figure 2 illustrates c-full cells and adjacency rules in the case $d^1 = 3, d^2 = 2$. As corresponding flag of faces, we have chosen $\{11\} \subset \{11, 21\} \subset \{11, 21, 31\} \subset \{11, 21, 31, 12, 22, 32\}$. The resulting path visits F_2 (the bottom face of the simplotope), but returns to F_1 . The second time F_2 is entered results in finding a c-full cell in F_2 (i.e., c-full relative to the bottom face), from which we enter F_3 (the whole simplotope, see right part of Figure 2). Finally a cell labelled 11, 21, 31, 32 is reached, i.e., this cell is a c-full one.

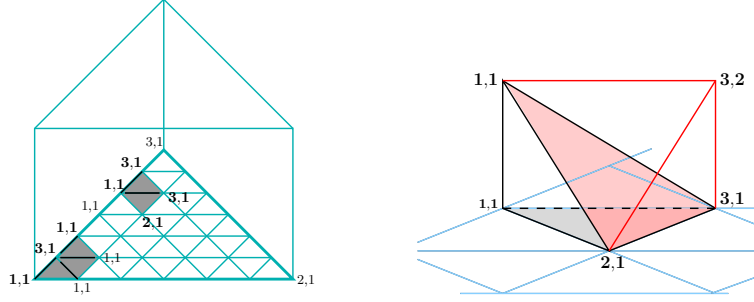


Figure 2: Locating c-full cells.

3 Nash equilibria

In this section, we will apply the simplotopal Sperner's Lemma to non-cooperative n -person games. A *non-cooperative n -person game* is a tuple $\Gamma = (S^1, \dots, S^n, u^1, \dots, u^n)$ consisting of strategy sets $S^i \subseteq \mathbb{R}^{d^i}$ and payoff functions $u^i : \prod_{i=1}^n S^i \rightarrow \mathbb{R}$. Denote by $N := \{1, \dots, n\}$ the set of players. Define $S := \prod_{i \in N} S^i$ and for $j \in N$ we write $S^{-j} := \prod_{i \in N, i \neq j} S^i$. We will make use of the following assumptions:

Each set S^i ($i = 1, \dots, n$) is the standard simplex in dimension d^i , so that S is a simplotope. Payoff functions are assumed to be continuous and $u^i(\bullet, s^{-i}) : S^i \rightarrow \mathbb{R}$ should be concave for each $i \in N$ and $s^{-i} \in S^{-i}$.

A *Nash equilibrium* in Γ is a strategy profile $\bar{s} = (\bar{s}^1, \dots, \bar{s}^n)$ in S such that for all $i \in N$ and all $s^i \in S^i$ we obtain $u^i(\bar{s}) \geq u^i(s^i, \bar{s}^{-i})$. Direct proofs for general existence of Nash equilibria involve Kakutani's fixed point theorem for correspondences. However, Geanakoplos (2003) gives a proof for existence that uses the Brouwer fixed point theorem for functions; therein, he replaces the upper-semicontinuous best-response correspondence by a “satisficing” function. As this function is “close” to the correspondence, its fixed points are close to those of the correspondence. There for each $i \in N$ a function $\varphi^i : S \rightarrow S^i$ is defined by

$$(1) \quad \varphi^i(\bar{s}) := \operatorname{argmax}_{s^i \in S^i} \{u^i(s^i, \bar{s}^{-i}) - \|s^i - \bar{s}^i\|^2\}.$$

As Geanakoplos (2003) shows, $\varphi^i(\bar{s})$ is not necessarily a best response of player i to \bar{s}^{-i} , but it “moves in the direction” of a best response. The quadratic term ensures uniqueness of the maximizer and thus φ is well defined

by (1). If \bar{s}^i is already a best response to \bar{s}^{-i} , then we obtain $\bar{s}^i = \varphi^i(\bar{s})$ and any fixed point of $\varphi = (\varphi^1, \dots, \varphi^n)$ is an equilibrium of Γ .

We use the functions φ^i to define a strategy-labelling ℓ on the simplotope S in the following way: For given $\bar{s} \in S$ calculate for each $i \in N$ the vector $\varphi^i - \bar{s}^i$ and assign a minimal coordinate as the label of player i at \bar{s} , i.e.,

$$(2) \quad \ell^i(\bar{s}) := \min \left\{ k \in L^i \mid \bar{s}_k^i > 0, k \in \operatorname{argmin}_{j \in L^i} (\varphi_j^i(\bar{s}) - \bar{s}_j^i) \right\}$$

for all $i \in N$. Set the label of $\bar{s} \in S$ as $\ell(\bar{s}) := (\ell^1(\bar{s}), \dots, \ell^n(\bar{s}))$.

Lemma 3.1.

The strategy-labelling ℓ is well-defined and is a simplotopal Sperner labelling on the set of strategy profiles S .

Proof. First we show that the set in (2) over which the minimum is taken is not empty. Suppose to the contrary that for $\bar{s}^i \in S^i$ and every k in the set $\operatorname{argmin}_{j \in L^i} (\varphi_j^i(\bar{s}) - \bar{s}_j^i)$ we have $\bar{s}_k^i = 0$. Hence $\varphi_k^i(\bar{s}) - \bar{s}_k^i \geq 0$ and as the sum over all coordinates is zero, it follows $\varphi^i(\bar{s}) - \bar{s}^i = 0$, so that $\operatorname{argmin}_{j \in L^i} (\varphi_j^i(\bar{s}) - \bar{s}_j^i) = L^i$ holds. Yet, this is a contradiction to $\bar{s}_k^i = 0$ for all minimizers.

From the above argument and the construction in (2), we can immediately conclude $\bar{s}_{\ell^i(\bar{s})}^i > 0$, which a simplotopal Sperner labelling is supposed to satisfy. \square

Now, let T be a triangulation of the simplotope. From Theorem 2.1 we know that T contains a c-full cell σ , whose label projections $L_i(\sigma)$ achieve the entire label set L^i , for each i .

As we shall see in the next two theorems, a c-full cell in S is important because: (1) it is approximate fixed point of the function φ , which is an equilibrium of Γ , and (2) sequences of c-full cells for smaller and smaller triangulation mesh sizes can converge to a fixed point of φ .

The following rather standard theorem is of importance in showing that full cells are approximate fixed points, when something about the function is known.

Theorem 3.2. *Given ϵ , let δ be the quantity associated with the absolute continuity of the function φ . Let ℓ be the labelling associated to φ defined in (2). Let σ_δ be a c-full cell of the triangulation (T, ℓ) of mesh size δ and let x be the barycenter of σ_δ . Then $|\varphi(x) - x| \leq 2n\epsilon$.*

Proof. Given a c-full cell σ_δ , let x be the barycenter of this cell, and suppose that all its vertices are located at $x + \bar{\delta}_i$. Then $|\bar{\delta}_i| < \delta$. Let $D_i(x) = \varphi_i - x_i$, the difference of $\varphi(x)$ and x in the coordinate x_i . By absolute continuity,

$$-\epsilon \leq D_i(x + \bar{\delta}_i) - D_i(x) \leq \epsilon.$$

Summing over all i , and using the fact that $\sum_i D_i(x) = 0$, we have

$$-n\epsilon \leq \sum_i D_i(x + \bar{\delta}_i) \leq n\epsilon.$$

In fact all the terms of the sum are negative because σ was a c-full cell. So we have

$$\sum_i |D_i(x + \bar{\delta}_i)| \leq n\epsilon.$$

But we also have that $|D_i(x)| \leq |D_i(x + \bar{\delta}_i)| + \epsilon$. So

$$|\varphi(x) - x| \leq \sum_i |D_i(x)| \leq \sum_i |D_i(x + \bar{\delta}_i) + \epsilon| \leq 2n\epsilon.$$

□

The above theorem gives a bound on how far a c-full cell is from a fixed point of φ . The next theorem shows that convergent sequences of c-full cells will converge to a fixed point of φ .

Theorem 3.3. *Let σ_δ be a c-full cell of the triangulation (T, ℓ) of mesh size δ and let s_δ be in σ_δ . As δ goes to zero, any convergent subsequence of s_δ converges to a fixed point of φ .*

Proof. We need to show that s_δ converges to a fixed point of $\varphi = (\varphi^1, \dots, \varphi^n)$. Let for all $i \in N$ and $j \in L^i$ be $v_\delta^{i,j}$ be a vertex of σ_δ , player i 's label of which is j . We can always find such vertices, since σ_δ is a c-full cell. Assume w.l.o.g. that s_δ converges to $\bar{s} \in S$ and hence $v_\delta^{i,j} \rightarrow \bar{s}$ holds for all i, j .

Because φ^i is continuous, we also conclude $\varphi^i(v_\delta^{i,j}) \rightarrow \varphi^i(\bar{s})$. Now, for $i \in N$ and $j \in L^i$ we have $(\varphi^i(v_\delta^{i,j}) - v_\delta^{i,j})_j \leq 0$ for all δ . Hence, in the limit we get $(\varphi^i(\bar{s}) - \bar{s})_j \leq 0$. Yet, the coordinate sum has to be zero, and therefore $\varphi^i(\bar{s}) = \bar{s}$. \square

So, the problem of finding an approximate Nash equilibrium is reduced to the problem of finding a c-full cell in the simplotope. This may be achieved by employing a path-following algorithm as described in the previous section (by choosing a flag of faces of the simplotope).

We note that the procedures of van der Laan & Talman (1982) and Freund (1986) would behave quite differently in this context. First of all, the resulting labelling of vertices in the triangulation would not be a vector of strategies, but a pair (j, k) where j is a player and k is a choice of strategy. The set of all such (j, k) is in one-to-one correspondence with the facets of S . The labelling suggested by van der Laan & Talman (1982) would label a vertex by asking for the strategy that decreases the most under φ^i , not just over one player's strategies, but over all players. This requires an interpersonal comparison of utilities that may be difficult to demand or exact from players. By contrast, our labelling function disentangles these player preferences into their own coordinates in the labelling.

We close this section with a remark on a work by Becker & Chakrabarti (2005). There the authors show the existence of Nash equilibria using a "better response map" and the BFPT and compare their technique to Geanakoplos (2003), by saying that "he approximates the expected utility function of a player by a strictly concave function to elicit a best reply function rather than a correspondence. Our technique applies to quasi-concave payoff functions in each players' strategy vector without resorting to a perturbation argument."

More precisely, for $s \in S$ and $x^i \in S^i$, let $\phi_i(s, x^i) = \max\{0, u_i(s_{-i}, x^i) - u_i(s)\}$. Each u^i is assumed continuous, and S^i compact, convex set in Euclidean space. For each s^{-i} , $u_i(s^{-i}, \cdot)$ is quasi-concave. Define $b_i : S \rightarrow S^i$ by

$$b_i(s) = \frac{s_i + \int_{S^i} \phi_i(s, a) a \, da}{1 + \int_{S^i} \phi_i(s, a) da}.$$

The mapping $b = (b^1, \dots, b^n)$ extends Nash's original idea of a "better response" in the existence proof for bimatrix games. Becker & Chakrabarti

(2005, Thm. 2) show that b has a fixed point and any fixed point is a Nash equilibrium. Due to this fact, we may as well construct a simplotopal Sperner labelling as in (2) using the function b instead of φ and finally arrive at an approximate Nash equilibrium by means of a path following argument.

4 Examples and Discussion

We now indicate how our results can be used to compute (approximate) Nash equilibria in two examples. In the second example, we compare our algorithm to that of Lemke & Howson (1964) for a bimatrix game where the the payoff functions are linear. We note, however, that unlike the Lemke & Howson algorithm, our algorithm works even when the payoff functions are not linear.

Example 4.1 (Cournot-Nash Duopoly).

We consider a Cournot model with two firms that can set supply quantities for a single good. The price for the good is determined by the function $p : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ with

$$p(q) := \max(0, \alpha - \beta q) \quad (\alpha, \beta > 0),$$

where q denotes the total supply. Production costs of firm i are given by $c^i : \mathbb{R}_+ \rightarrow \mathbb{R}$ with $c^i(q) := \gamma^i q$ ($i = 1, 2$) with constant positive marginal costs γ^i . To obtain positive equilibrium quantities, we will assume $\gamma^i < \alpha$ ($i = 1, 2$). Clearly, no firm will be willing to produce $q^i > \alpha/\beta$, which would result in a price equal to zero. Hence we may assume, that each firm's strategy set is given by $S^i = [0, \alpha/\beta]$. In view of our model, we will refer to the two extreme quantities as the pure strategies a firm possesses. Any supply quantity can be expressed as a convex combination of these two pure strategies. Firms' payoffs depend on both quantity choices and are given by $u^i(q^1, q^2) := q^i p(q^1 + q^2) - c^i(q)$. Best responses are single valued and given by functions $\Phi^i : S^{3-i} \rightarrow S^i$ with $\Phi^i(q^{3-1}) = \max\left\{0, \frac{\alpha - \beta q^{3-i} - \gamma^i}{2\beta}\right\}$ ($i = 1, 2$).

To determine the labelling function on S , we make the following observation: Fix $q^2 \in S^2$ and take $q^1 > \Phi^1(q^2)$. Due to concavity assumptions on payoff functions, we know, that $u^1(q^1, q^2) > u^1(q^1 + r, q^2)$ is true for positive r . Hence we may conclude $\varphi^1(q^1, q^2) < q^1$ is satisfied. Analogous arguments hold for

$q^1 < \Phi^1(q^2)$. Therefore we get $\varphi^1(q^1, q^2) - q^1 < 0$, if and only if $q^1 > \Phi^1(q^2)$. Remembering that the first pure strategy corresponds to *zero supply* and the second to *maximal supply*, we have $\ell^i(q^1, q^2) = 1$, if $q^i < \Phi^i(q^{3-i})$ and $\ell^i(q^1, q^2) = 2$, if $q^i > \Phi^i(q^{3-i})$.

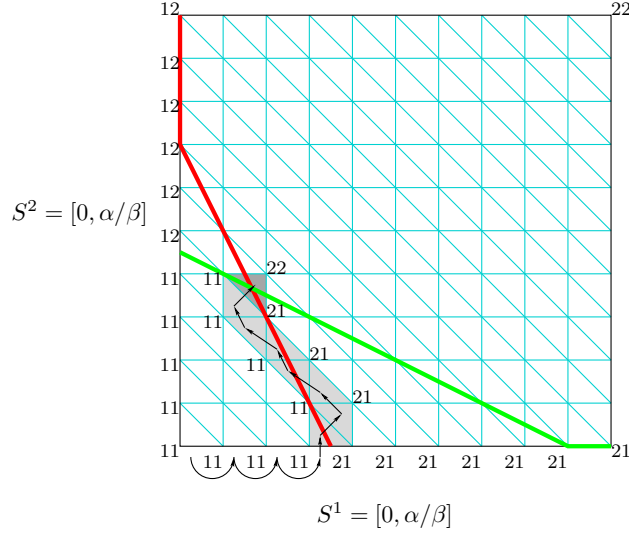


Figure 3: Cournot Duopoly labelling for $\gamma^1 = \frac{3}{10}\alpha, \gamma^2 = \frac{1}{10}\alpha$

Figure 3 shows best responses for the two firms as well as a correspondingly labelled triangulation. If we start the algorithm with the flag

$$\begin{aligned} F_0 &:= \{(e^1, e^1)\} \subset F_1 := \text{conv}\{(e^1, e^1), (e^2, e^1)\} \\ &\subset F_2 := \text{conv}\{(e^1, e^1), (e^1, e^2), (e^2, e^1), (e^2, e^2)\} = S \end{aligned}$$

(which we abbreviate by corresponding labels, i.e. $\{\{11\}, \{11, 21\}, \{11, 21, 12, 22\}\}$), then we follow the path indicated in the figure and obtain a c-full cell. In the case $\gamma^1 = \frac{3}{10}\alpha, \gamma^2 = \frac{1}{10}\alpha$, which is depicted, the c-full cell has vertices $\frac{\alpha}{\beta}(\frac{1}{10}, \frac{4}{10}), \frac{\alpha}{\beta}(\frac{2}{10}, \frac{4}{10}), \frac{\alpha}{\beta}(\frac{2}{10}, \frac{3}{10})$. Supply quantities in equilibrium are $\bar{q}^i = \max(0, \frac{\alpha - 2\gamma^i + \gamma^{3-i}}{3\beta})$, which means $\bar{q}^1 = \frac{1}{6}\frac{\alpha}{\beta}, \bar{q}^2 = \frac{11}{30}\frac{\alpha}{\beta}$ for the case in Figure 3. Hence, the equilibrium point is located in the c-full cell determined by the algorithm. \square

Example 4.2.

We consider the bimatrix game $\Gamma = (\mathbb{X}, \mathbb{Y}, A, B)$ with strategy sets

$\mathbb{X} := \{x \in \mathbb{R}_+^3 \mid \sum_{i=1}^3 x_i = 1\}$, $\mathbb{Y} := \{y \in \mathbb{R}_+^2 \mid \sum_{i=1,2} y_i = 1\}$. Payoff functions are determined by 3×2 matrices A and B so that $u^1(x, y) = x^t A y$, $u^2(x, y) = x^t B y$ holds. We set

$$A := \begin{pmatrix} 2 & 0 \\ 1 & 2 \\ 0 & 3 \end{pmatrix} \quad B := \begin{pmatrix} 3 & 0 \\ 0 & 2 \\ 2 & 1 \end{pmatrix}.$$

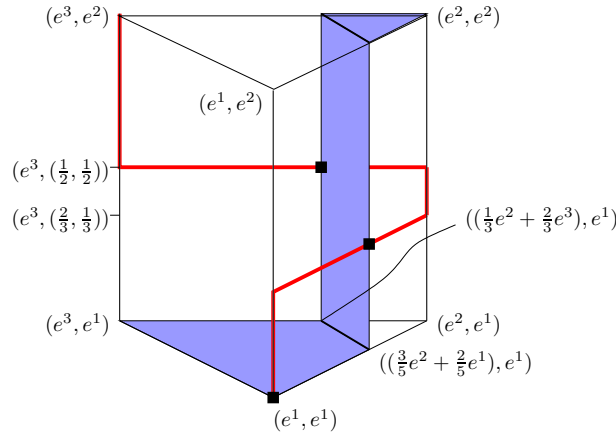


Figure 4: Best reply correspondences in Γ

Figure 4 depicts both best reply correspondences in $S := \mathbb{X} \times \mathbb{Y}$. The shaded area are player 2's best replies, whereas the thick line represents player 1's best responses. The three intersection points are the three Nash equilibria of Γ . They are given by $\bar{x} := (e^1, e^1)$, $\tilde{x} := ((\frac{2}{5}, \frac{3}{5}, 0), (\frac{2}{3}, \frac{1}{3}))$ and $\hat{x} := ((0, \frac{1}{3}, \frac{2}{3}), (\frac{1}{2}, \frac{1}{2}))$.

The labelling function ℓ as a function on S is given as follows. Player 2 has only two pure strategies, so that, as in the previous example, ℓ^2 is given by her best reply correspondence. To each point (x, y) on the graph of player 2's best reply correspondence, we can assign either label $\ell^2(x, y) = 1$ or $\ell^2(x, y) = 2$ (obey the boundary conditions for $y = e^1$ or $y = e^2$!). The graph subdivides S into two regions. Points in the region that contains the vertex (e^1, e^2) are labelled 2, whereas the remaining points are assigned label 1.

Next, we determine ℓ^1 . It turns out that for matrix games the function φ^i

reduces to

$$\varphi^i(x, y) = \operatorname{argmin} \left\{ \|x' - (\tfrac{1}{2}Ay)\|^2 \mid x' \in \mathbb{X} \right\}.$$

(see Geanakoplos (2003) or Lemke & Howson (1964))¹, i.e. we compute the point in \mathbb{X} , which is closest to the vector $\frac{1}{2}Ay \in \mathbb{R}^n$. Figure 5 shows how ℓ^1 looks like for different $y \in \mathbb{Y}$. One readily verifies that we indeed have a Sperner labelling of S .

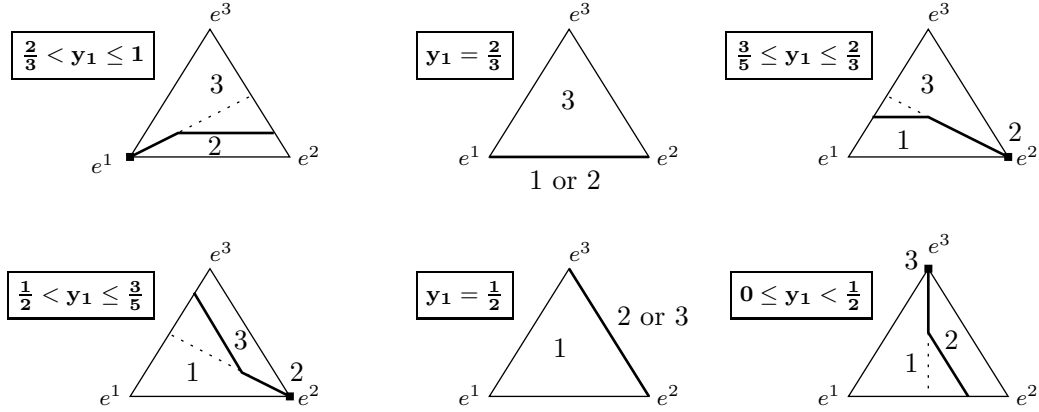


Figure 5: Labeled slices of \mathbb{X}

We will now briefly describe, which path is followed, when we choose (e^3, e^2) as starting vertex and pick the flag

$$\begin{aligned} F_0 &:= \{(e^3, e^2)\} \subset F_1 := \operatorname{conv}((e^3, e^2), (e^2, e^2)) \\ &\subset F_2 := \operatorname{conv}((e^3, e^2), (e^2, e^2), (e^3, e^1), (e^3, e^2)) \subset F_3 := S. \end{aligned}$$

The starting point is labelled $\{32\}$. The path in the procedure follows face F_1 until we hit a cell labelled $\{32, 22\} = L(F_1)$. From the bottom right picture in Figure 5, we can see that this is already fulfilled by the starting cell. Next, according to the chosen flag, we enter $F_2 = \mathbb{X} \times \{e^2\}$ and follow faces labelled $\{32, 22\}$ until we hit a cell, which is c-full in F_2 (i.e. labelled $\{32, 22, 12\}$). Again, this is immediately the case, in view of Figure 5. Now entering $F_3 = S$, we proceed to move from one simplex to the unique neighboring simplex that shares a face with labels $\{32, 22, 12\}$, until we come to a cell σ with $L^2(\sigma) = \{1, 2\}$. This involves following a path through cells that “stick to $e^3 \in \mathbb{X}$ ” and moves “towards $e^1 \in \mathbb{Y}$ ” until the second coordinate is

¹The reduction in the bimatrix case involves the Kuhn-Tucker theorem.

approximately at $\frac{1}{2}e^1 + \frac{1}{2}e^2$ (it cannot go further down, since cells below have no face with label 12). The path now directly finds the equilibrium \hat{x} . In Figure 6 the path to the equilibrium is illustrated. The algorithm visits cells along the solid line starting at (e^3, e^2) and ending in \hat{x} . \square

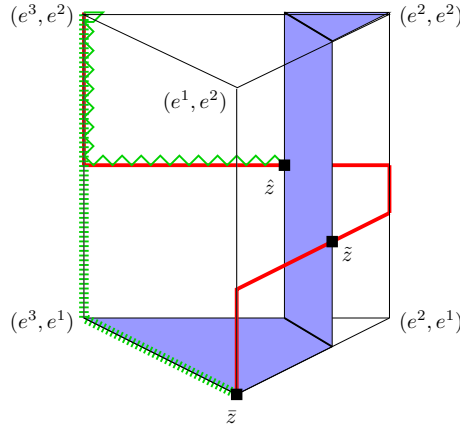


Figure 6: Paths to equilibria

Observe that the final full cell that the algorithm determines in Example 4.2 depends on the whole flag chosen. For the same starting point (i.e., the first face in the flag) as above, but flag (indicated by the vertex labels) $\{32\}, \{32, 31\}, \{32, 31, 11, 12\}, \{32, 31, 11, 12, 21, 22\}$ the corresponding path “runs down” in S and then finds \bar{x} as equilibrium (see the dashed path in Figure 6).

Comparing our algorithm to the one introduced by Lemke & Howson (1964) (hereafter LH) for bimatrix games, which was later extended by Rosenmüller (1971) to the n -person case, the two algorithms appear to follow similar paths. The LH algorithm winds its way along edges in the simplotope S . The path is determined by an “almost equilibrium condition”, meaning that at most one player playing at most one pure strategy with positive probability, which is not a best response to the opponent’s strategy. This uniquely determines a path from a vertex of S to an equilibrium. Thereby, the players alternately adjust their strategies, one in each step. Considering Example 4.2, the vertex (e^3, e^2) is a legal starting point for the LH algorithm, which finds in this case also the equilibrium \hat{x} , taking (approximately) the same way as we described above. Yet, the LH algorithm requires a first

move of player 2. In our model this corresponds to choosing a flag with $F_1 := \{32, 31\}$, which results in a different path than the LH algorithm takes. By appropriate choice of F_2 , our algorithm can even get to a different equilibrium, namely \bar{x} . This shows that, even if we try to mimic the LH algorithm, we may end up with a different path and even a different equilibrium point.

Another difference to the LH algorithm is that the path we follow may involve simultaneous movements of the players to find an equilibrium.² This is because our algorithm follows a different rationale to determine an (approximate) equilibrium. Each face in the flag is the product of simplices spanned by specific pure strategies of both players. For example, the face with labels $\{32, 31, 21, 22\}$ includes all strategies, in which player 1 randomizes only over his second and third pure strategy, whereas player 2 is allowed to pick any mixed strategy. The idea of our algorithm is the following: As we enter a new face in the flag, we allow one of the players to include an additional pure strategy. Finding a full cell relative to face F_k means determining an approximate equilibrium of the game, in which only those pure strategies are available, that actually span the face F_k . This means, during the course of the algorithm we consider restrictions of the game³ and determine one of its equilibria. Entering a new face of the flag, we include more information of the game. Moving in the last face (which is S itself) means taking the complete game into account. Clearly, movements to find a c-full cell in face F_k can involve adjustments of both players (mixed) strategies.

To see this, note that a c-full cell σ_k relative to F_k is one face of a cell σ_{k+1} in F_{k+1} . Roughly spoken, we enter the face F_{k+1} through σ_k and continue by leaving σ_{k+1} through a face with the same set of labels as σ_k has.

Like the LH-algorithm, our algorithm does not find the equilibrium \tilde{x} in Example 4.2 for any possible choice of flags. However, the class of games is not restricted to those with multilinear payoff functions.

Finally, we want to comment on an extensions of the class of games that can be considered with our approach. Geanakoplos (2003) proved existence of Nash equilibria for games with convex compact strategy sets and, as above, concave payoff functions by using the function φ . With similar techniques developed for Theorem 2.1, one might at least achieve a similar result for

²Take the flag $\{\{32\}, \{32, 31\}, \{32, 31, 21, 22\}, \{32, 31, 21, 22, 11, 12\}\}$.

³i.e., restrictions on available strategies.

polytopal strategy spaces. The algorithm then operates in the product of polytopes. With appropriate modifications of flags and rules for moving from cell to cell, one reaches an approximate equilibrium in the same fashion.

We close with two open questions. As our algorithm winds its way through simplices of a given triangulation, it does not describe a *precise* path through S . However, one may obtain an *exact* version by following points in S that potentially can carry all labels of a corresponding c-full cell. But, a decomposition of S into regions that are labelled the same way involves non-convex sets (see, e.g., Figure 5). A second question concerns the parity of the number of c-full cells. For example, a non-degenerate⁴ bimatrix game exhibits an odd number of Nash equilibria. As De Loera, Peterson & Su (2002) show, there is an odd number of c-full cells for each choice of a flag. Since a c-full cell may appear for different flags, one cannot conclude that the overall number of c-full cells is odd and thus reach a similar conclusion in the bimatrix case.

References

- BECKER, R. A. & S. K. CHAKRABARTI (2005): “Satisficing behavior, Brouwer’s fixed point theorem and Nash equilibrium”, *Economic Theory*, 26(1), 63–83.
- BORDER, K. C. (1985): *Fixed Point Theorems with Applications to Economics and Game Theory*. Cambridge University Press, Cambridge, UK.
- DE LOERA, J. A., E. PETERSON & F. E. SU (2002): “A polytopal generalization of Sperner’s lemma”, *J. Combin. Theory Ser. A*, 100(1), 1–26.
- FREUND, R. M. (1986): “Combinatorial theorems on the simplotope that generalize results on the simplex and cube”, *Mathematics of Operations Research*, 11(1), 169–179.
- GEANAKOPOLOS, J. (2003): “Nash and Walras equilibrium via Brouwer”, *Economic Theory*, 21, 585–603.

⁴see Lemke & Howson (1964) for a definition.

- HERINGS, P. & A. H. VAN DEN ELZEN (2002): “Computation of the Nash equilibrium selected by the tracing procedure in n -person games”, *Games and Economic Behavior*, 38, 89–117.
- LEMKE, C. & J. HOWSON (1964): “Equilibrium points in bimatrix games”, *SIAM Journal on Applied Mathematics*, 12.
- ROSENMÜLLER, J. (1971): “On a generalization of the Lemke-Howson algorithm to noncooperative n -person games”, *SIAM Journal on Applied Mathematics*, 21(1), 73–79.
- TODD, M. J. (1976): *The Computation of Fixed Points and Applications*. Springer-Verlag, Berlin, Lecture Notes in Economics and Mathematical Systems, Vol. 124.
- VAN DER LAAN, G. & A. J. J. TALMAN (1982): “On the computation of fixed points in the product space of unit simplices and an application to noncooperative n person games”, *Mathematics of Operations Research*, 7(1), 1–13.
- YANG, Z. (1999): *Computing Equilibria and Fixed Points : The Solution of Nonlinear Inequalities*, vol. 21 of *Theory and Decision Library. Series C: Game theory, mathematical programming, and mathematical economics*. Kluwer Academic Publishers, Boston, Mass.