

Geometry and equilibria in bimatrix games

Anne Verena Balthasar

A thesis submitted for the degree of
Doctor of Philosophy

Department of Mathematics
London School of Economics
and Political Science

September 2009

Declaration

I certify that the thesis I have presented for examination for the MPhil/PhD degree of the London School of Economics and Political Science is my own work, apart from Chapters 2 and 4, which are joint work with Bernhard von Stengel.

The copyright of this thesis rests with the author. Quotation from it is permitted, provided that full acknowledgement is made. This thesis may not be reproduced without the prior written consent of the author.

I warrant that this authorization does not, to the best of my belief, infringe the rights of any third party.

Abstract

This thesis studies the application of geometric concepts and methods in the analysis of strategic-form games, in particular bimatrix games. Our focus is on three geometric concepts: the index, geometric algorithms for the computation of Nash equilibria, and polytopes.

The contribution of this thesis consists of three parts. First, we present an algorithm for the computation of the index in degenerate bimatrix games. For this, we define a new concept, the “lex-index” of an extreme equilibrium, which is an extension of the standard index. The index of an equilibrium component is easily computable as the sum of the lex-indices of all extreme equilibria of that component.

Second, we give several new results on the linear tracing procedure, and its bimatrix game implementation, the van den Elzen-Talman (ET) algorithm. We compare the ET algorithm to two other algorithms: On the one hand, we show that the Lemke-Howson algorithm, the classic method for equilibrium computation in bimatrix games, and the ET algorithm differ substantially. On the other hand, we prove that the ET algorithm, or more generally, the linear tracing procedure, is a special case of the global Newton method, a geometric algorithm for the computation of equilibria in strategic-form games. As the main result of this part of the thesis, we show that there is a generic class of bimatrix games in which an equilibrium of positive index is not traceable by the ET algorithm. This result answers an open question regarding sustainability.

The last part of this thesis studies the index in symmetric games. We use a construction of polytopes to prove a new result on the symmetric index: A symmetric equilibrium has symmetric index $+1$ if and only if it is “potentially unique”, in the sense that there is an extended symmetric game, with additional strategies for the players, where the given symmetric equilibrium is unique.

Acknowledgements

First and foremost, I would like to thank my supervisor, Bernhard von Stengel, for his support and enthusiasm. Working with him has been an insightful and rewarding experience. Thanks for answering questions patiently, taking the time for endless whiteboard discussions and reading and improving my drafts, time and again.

I would like to thank Stephan, who gave me the idea to go to the UK in the first place, and Felix, for making me fall in love with London. In the LASE maths department, I am grateful to Jan, Joel, Graham, David, Marta and Julian for all those nice lunchtime discussions, and for making my time here so much more fun. I am thankful to Luis, Peter and Viresh who made my start at the LSE much easier. Thanks to Rahul and Arndt for answering lots of questions. Also, I would like to thank Jackie, Simon, Dave and Carl for being such a helpful admin team, and Vlad for resolving all sorts of computer issues.

I would like to thank the LSE research studentship scheme and the EPSRC for financial support.

A special thanks goes to all my friends, who made difficult times much more bearable, especially Iris and little Oskar. I would also like to thank Andrew and Shana for very helpful discussions.

And finally, the biggest thanks goes to Martin, for everything.

Contents

1	Introduction	7
1.1	Geometry and game theory	7
1.2	Thesis outline	10
1.3	Preliminaries	14
2	The index of an equilibrium component	21
2.1	Introduction	21
2.2	The index	23
2.3	A lexicographically perturbed game	31
2.4	Computing the index of an equilibrium component: The lex-index	37
3	Equilibrium Tracing in Strategic-Form Games	42
3.1	Introduction	42
3.2	Van den Elzen-Talman versus Lemke-Howson	44
3.3	Relationships to the global Newton method	46
3.4	Traceability and the index of equilibria	52
3.5	Open questions	60
4	Index and Uniqueness of Symmetric Equilibria	61
4.1	Introduction	61
4.2	The symmetric index	64
4.3	Polytopes and symmetric equilibria	66
4.4	Oriented facets	70
4.5	From unit vector to symmetric games	74
4.6	P-matrix prisms	77
4.7	Re-arranging the polytope P^Δ	83
4.8	Disjoint completely labelled facets	87

Contents

4.9	General completely labelled facets	94
4.10	Open questions	101
	Bibliography	102

1

Introduction

1.1 Geometry and game theory

Geometric ideas are ubiquitous in game theory. Many game theoretic problems lead naturally to geometric or topological questions. An important example is the notion of Nash equilibrium, which is the central solution concept in the theory of strategic-form games. The set of Nash equilibria of a given game carries a natural geometric structure: It is the set of solutions of a system of polynomial equations and inequalities, i.e. a semi-algebraic set. In the case of bimatrix games (i.e. two-player games in strategic form), this set has an even simpler structure: It is given by a set of linear constraints, coupled with a complementarity condition, hence is the solution set of a linear complementarity problem. As such, it is a finite union of polyhedra. This nice geometric structure has been exploited in various aspects of the study of Nash equilibria.

This thesis explores three geometric concepts that have been widely studied in the context of strategic-form (and especially bimatrix) games: Polytopes, geometric algorithms for the computation of Nash equilibria, and the index of an equilibrium. Before we describe the contribution of this thesis, we would like to give a short overview of these three concepts in the context of game theory, and explain how they are linked.

One of the most fundamental geometric concepts used in bimatrix game theory is that of polyhedra and polytopes. Polytopes have been used both to visualize

game theoretic ideas and to solve game theoretic problems. A very useful tool in the study of bimatrix games is the “best-reply polytope”, which is valuable in the study of Nash equilibria in various respects: Von Stengel (1999) used these polytopes to refute the conjecture by Quint and Shubik (1997) on the maximal number of Nash equilibria that a $d \times d$ bimatrix game can have. Von Schemde (2005) studied stability properties of equilibria using a construction based on best-reply polytopes. Savani and von Stengel (2006) applied best-reply polytopes to answer a long-standing open question in algorithmic game theory, namely if the classic algorithm for the computation of Nash equilibria, the Lemke-Howson algorithm, has exponential running time.

The computation of one or all Nash equilibria of a bimatrix game is closely linked to polytopal concepts. For this reason, polytopes constitute a central tool in algorithmic game theory. For example, the extreme equilibria of a bimatrix game correspond to certain vertices of the best-reply polytope. Hence computing all Nash equilibria of a bimatrix game essentially corresponds to vertex enumeration (Vorob’ev, 1958; Avis et al., 2010). Likewise, polytopes are used in the study of complementary pivoting algorithms for the computation of a single equilibrium of a bimatrix game. As mentioned earlier, the set of Nash equilibria in bimatrix games has a particularly nice geometric structure as the solution set of a linear complementarity problem. A general algorithm for the solution of such a linear complementarity problem is Lemke’s algorithm (Lemke, 1965), which is a complementary pivoting method that walks along edges of a suitable polyhedron. Two of the best-known algorithms for the computation of equilibria in bimatrix games, those by Lemke and Howson (1964) and van den Elzen and Talman (1991), are both special cases of Lemke’s algorithm (Savani, 2006; von Stengel et al., 2002), and as such have a straightforward geometric interpretation in terms of polyhedra.

However, the use of geometric tools for equilibrium computation has not been restricted to bimatrix games. Geometric algorithms have also been developed for the computation of Nash equilibria in general strategic-form games. An example is the global Newton method by Govindan and Wilson (2003a), which is based on the particularly nice geometry of the graph of the equilibrium correspondence. The equilibrium correspondence on the space of strategic-form games of a given dimension maps each game to its set of equilibria. The graph of this correspon-

dence is a topological manifold, whose one-point-compactification is homeomorphic to a sphere (Kohlberg and Mertens, 1986). This simple geometric structure is exploited in the global Newton method for the construction of a geometric algorithm that computes a Nash equilibrium of a given strategic-form game. This algorithm is an example of the wider class of homotopy (or path-following) algorithms, which trace a path in a generically one-dimensional manifold in order to find an equilibrium (Herings and Peeters, 2010). Another example of such a homotopy algorithm is the linear tracing procedure introduced by Harsanyi (1975), which generalizes the van den Elzen-Talman algorithm from bimatrix to general strategic-form games. The tracing procedure plays a crucial role in the equilibrium selection theory developed by Harsanyi and Selten (1988), and in an equilibrium refinement concept suggested by Myerson (1997).

Algorithmic problems in game theory are closely related to another important geometric concept, the index. The index of an equilibrium is a topological notion which assigns to each connected component of Nash equilibria an integer, such that the indices of all equilibrium components of a game add up to one. The index was developed by Shapley (1974) in the context of the Lemke-Howson algorithm, and all algorithms mentioned above have in common that generically, they will only find equilibria of positive index (Garcia and Zangwill, 1981).

But the relevance of the index has grown way beyond algorithmic issues. The index of an equilibrium component carries crucial information about many of its properties. For this reason, the index plays a considerable role in equilibrium refinement and selection theory. In nondegenerate bimatrix games, equilibrium components consist of isolated points, whose index can be either $+1$ or -1 . In such games, it has been shown that several important properties of an equilibrium depend on its index. As already mentioned, homotopy algorithms for the computation of equilibria, like the Lemke-Howson or van den Elzen-Talman algorithms, will generically find only equilibria of index $+1$ (Garcia and Zangwill, 1981). Von Schemde (2005) proved that an equilibrium has positive index if and only if it can be made the unique equilibrium of an extended game, where strategies with suitable payoffs are added. Furthermore, the index carries crucial information about the dynamic stability of an equilibrium with regards to Nash fields (i.e. vector fields that have exactly the equilibria of the given game as rest points). Positively indexed equi-

libria can be made dynamically stable by a suitable choice of Nash field, whereas equilibria with negative index are always unstable (Hofbauer, 2003).

In degenerate games, the index is no longer restricted to the values $+1$ or -1 but can take any integer as value (Govindan et al., 2003). In this case, the index of an equilibrium component carries information about the “essentiality” or “stability” of the component, in terms of payoff perturbations (as opposed to the “dynamic stability” considered earlier). Various concepts of stability have been suggested in the search for a satisfying theory of equilibrium refinement and selection, following the main idea that an equilibrium component is stable if it does not vanish under slight perturbations of the payoffs (see Kohlberg and Mertens (1986) and the subsequent literature). The concept that has been shown to best capture the interdependence between index and stability is that of hyperstability. An equilibrium component is called hyperstable if it is stable in every equivalent game, where two games are called equivalent if they can be reduced to the same game by deleting “superfluous” strategies that are convex combinations of other strategies. An equilibrium component is called uniformly hyperstable if the hyperstability condition holds uniformly over all equivalent games. Govindan and Wilson (2005) proved that an equilibrium component is uniformly hyperstable if and only if it has nonzero index.

We can conclude that the index, a purely geometric notion, is relevant to the study of equilibrium properties both in degenerate and nondegenerate bimatrix games. But geometric ideas are even more ubiquitous in game theory than our exposition suggests. To give an example, Mertens (1989, 1991) uses homology theory, a tool from algebraic topology, to study strategic stability. However, we restrict the scope of this thesis to the three geometric concepts outlined above: Polytopes, geometric algorithms for equilibrium computation, and the index. The main focus of this thesis is on bimatrix games, with a notable exception in our study of game theoretic algorithms in Chapter 3.

1.2 Thesis outline

The contribution of this thesis consists of three parts: As a first result, we present an algorithm to compute the index in degenerate bimatrix games. In the second part,

we give several new results on the linear tracing procedure of Harsanyi (1975), and its bimatrix game implementation, the van den Elzen-Talman algorithm. The last part studies the index in symmetric games; we use a construction of polytopes to prove a new result on the symmetric index.

In **Chapter 2**, we first give an exposition of the index of an equilibrium component in bimatrix games. Then, as the main result of this chapter, we present an algorithm for the computation of the index in degenerate games.

In nondegenerate games, the index is easily computable, essentially as the sign of a suitable determinant. However, in degenerate games, where equilibrium components can occur, there is no such straightforward method. Existing algorithms rely on perturbations of the payoffs of the game, or on interior approximations of a Nash field. In order to arrive at a simpler algorithm, we extend the definition of the index of isolated equilibria in nondegenerate games to extreme equilibria in degenerate games. We call this new index notion the *lex-index* of an extreme equilibrium.

The crucial ingredient for our algorithm is the following, intuitively appealing result: The index of an equilibrium component is the sum of the lex-indices over all extreme equilibria of that component. The lex-index of an extreme equilibrium is easily computable, using just the game matrices (A, B) , without resorting to topological concepts such as perturbations or interior approximations. Hence our method offers an improvement on existing algorithms.

This chapter is joint work with Bernhard von Stengel, intended for publication.

In **Chapter 3**, we analyze several geometric algorithms for the computation of Nash equilibria. Our focus is on the van den Elzen-Talman algorithm, a complementary pivoting method for equilibrium computation in bimatrix games. The algorithm starts at an arbitrary strategy profile, called prior. Both players adjust their strategies until an equilibrium is reached. This algorithm has the advantage of being more flexible than the classic algorithm for equilibrium computation in bimatrix games, the Lemke-Howson method: While the Lemke-Howson algorithm relies on a finite set of starting points, the van den Elzen-Talman algorithm can start anywhere in the strategy space. Another useful property of the latter algorithm is that it implements Harsanyi's and Selten's linear tracing procedure, which plays an

important role in equilibrium selection and refinement theory.

We answer several questions regarding these algorithms. First, we show that the Lemke-Howson and van den Elzen-Talman algorithms differ substantially: The van den Elzen-Talman algorithm, when started from a pure strategy and its best response as a prior, in general finds a different equilibrium than the corresponding Lemke-Howson method. Secondly, we prove that the van den Elzen-Talman algorithm, or more generally, the linear tracing procedure, is a special case of the global Newton method, a geometric algorithm for the computation of equilibria in general strategic-form games introduced by Govindan and Wilson (2003a). As the third and main result of this chapter, we show that the van den Elzen-Talman algorithm is not flexible enough to find every equilibrium of positive index. Our result is based on the concept of “traceability”: An equilibrium is called traceable if it is found by the van den Elzen-Talman algorithm from an open set of priors (Hofbauer, 2003). We prove that there is a generic class of bimatrix games in which an equilibrium of positive index is not traceable. This result answers an open question that arises from a closely related notion of sustainability: Myerson (1997) suggested to call an equilibrium sustainable if it is found by the tracing procedure from an open set of priors. Our result shows that in this sense, not all equilibria of positive index are sustainable.

A version of Chapter 3 has been published in *Economic Theory* (Balthasar, 2010).

In **Chapter 4**, we analyze the index in the context of symmetric bimatrix games. A bimatrix game is called symmetric if the players have the same number of strategies, and the two players are interchangeable. More precisely, this means that the payoff matrix of one player is the transpose of the payoff matrix of the other player. Symmetric games play an important role in evolutionary game theory, where a mixed strategy can represent the frequencies of individual pure strategies that occur in a population. A symmetric game may have both symmetric and non-symmetric equilibria. In certain situations – for example if the players have no way of determining which of the two possible player positions they are in – it makes sense to only consider the symmetric equilibria. In a symmetric game, the “symmetric index” of a symmetric equilibrium is defined analogously to the index in a general bimatrix game. For any symmetric equilibrium, its symmetric index may differ from the “usual” (i.e. bimatrix game) index, as can be seen in simple examples like

the game of chicken.

We prove that in a nondegenerate symmetric game, a symmetric equilibrium has symmetric index $+1$ if and only if it is “potentially unique” in the sense that there is an extended symmetric game, with additional strategies for the players, where the given symmetric equilibrium is unique. The corresponding statement for bimatrix games has been proved by von Schemde (2005). However, the symmetric case does not follow from the seemingly more general result on bimatrix games for the following reasons: First, as explained above, the bimatrix index and the symmetric index of a fixed symmetric equilibrium may differ. Secondly, the game needs to be extended in a symmetric way, but the extension in the corresponding result on bimatrix game is always asymmetric.

Our proof relies on a construction of polytopes, which should be of independent geometric interest. Nondegenerate symmetric $d \times d$ games correspond to simplicial d -polytopes whose vertices are labelled with labels from the set $\{1, \dots, d\}$. The symmetric equilibria correspond to completely labelled facets of that polytope, i.e. facets whose vertices have all labels in $\{1, \dots, d\}$ (apart from one completely labelled facet, which gives rise to an “artificial” equilibrium). Every completely labelled facet carries a natural orientation. To prove our result on the symmetric index, it suffices to prove the following statement in the corresponding polytopal setting: Whenever we have a pair of completely labelled facets of opposite orientation, we can add labelled points to the polytope such that the only completely labelled facets of the extended polytope are the two given ones. The proof of this polytopal result is based on ideas developed in von Schemde and von Stengel (2008), who use a very similar approach for a constructive proof of the corresponding result on the “usual” index of bimatrix games.

We derive the game theoretic result from its geometric counterpart by applying the above polytopal result to the polytope that corresponds to a given symmetric game. For a fixed symmetric equilibrium of positive index, we add points to this polytope such that the only two completely labelled facets of the extended polytope are the one which corresponds to the given equilibrium, and the “artificial” one. The added points are then used to define a suitable extension of the symmetric game. A central part in the step from added points to added strategies is played by a class of bimatrix games that we call *unit vector games*. These games generalize the

imitation games of McLennan and Tourky (2007), who prove that every symmetric game corresponds to a certain imitation game.

Starting from the added points in the extended polytope we create a symmetric extension of the given symmetric game in two steps: First, we use the added points to extend the corresponding imitation game to a unit vector game, by adding strategies for the column player. Secondly, we symmetrize this extended game by adding suitable payoff rows. This second part, the symmetrization, is the crucial step from the geometric to the game theoretic result. In the final symmetric extension, the given symmetric equilibrium is unique.

Chapter 4 is joint work with Bernhard von Stengel, intended for publication.

In the remainder of **this chapter**, we summarize relevant prerequisites and notational conventions which we use throughout this thesis.

1.3 Preliminaries

In this section we summarize some terminology and key results about games and polytopes. The contents of this section can be found in the standard literature on the subject. For strategic-form games, we refer the reader to Ritzberger (2002), for bimatrix games to von Stengel (2002, 2007). Details on the theory of polytopes can be found in Grünbaum (2003) or Ziegler (1995).

By vector we mean column vector (although for reasons of space, in examples we often write vectors as row vectors). Inequalities between vectors hold component-wise. As usual, e_i denotes the i th unit vector, and $\mathbf{0}$ and $\mathbf{1}$ the all-zero- and all-one-vector, respectively, with dimension understood from context. For a matrix C , we denote by C^\top its transpose. We write I for the identity matrix, and E for the matrix that has all entries equal to one. The dimension of these matrices may vary according to context. We write $C = [c_1 \cdots c_n]$ if C is a matrix with columns c_1, \dots, c_n . For a set X , we denote by $|X|$ its cardinality.

A (*finite*) *strategic-form game* is given by the following data: a finite set of players, and for each player, on the one hand a finite set of pure strategies that are available to him, and on the other hand his payoff function. A *strategy profile* is a tuple of strategies, one for each player. A player's payoff function assigns to each strategy

profile a real number, which is the payoff that this player receives from the strategy profile. Players are allowed to use *mixed strategies*, i.e. randomize over their pure strategies. The payoff function is extended to mixed strategies by taking the expected payoff. Denote by Δ_r the $r - 1$ -dimensional standard simplex:

$$\Delta_r = \{x \in \mathbb{R}^r \mid x^\top \mathbf{1} = 1, x \geq \mathbf{0}\}$$

The set of mixed strategies of a player who has r strategies can be identified with Δ_r , which we also refer to as *this player's strategy space*. We use the term *strategy space* (without reference to a specific player) for the set of all mixed strategy profiles, which is the cartesian product of the strategy spaces of all players. In a fixed “dimension” (i.e. when fixing the number of players, and for each player his number of strategies) a game is determined by the payoffs to the players from the pure strategy profiles, i.e. by a finite set of real numbers. The *space of games* in a fixed dimension is defined as the set of all games in that dimension. It can be identified with a suitably-dimensional real space \mathbb{R}^d .

The central solution concept for a strategic-form game is that of *Nash equilibrium*. A strategy profile is called a Nash equilibrium if no player has an incentive to deviate unilaterally, i.e. if no player can increase his payoff by changing his strategy while all other players adhere to the equilibrium profile. Nash (1951) proved that every strategic-form game has an equilibrium. The *equilibrium correspondence* is the set-valued function from the space of games to the strategy space that assigns to each game its set of equilibria.

A *bimatrix game* is a two-player strategic-form game, in which the payoffs are given by two $m \times n$ matrices (A, B) . The first player chooses a row as pure strategy, the second a column. The payoffs are then given by the respective matrix entry of A for the first player, and B for the second. We denote the set of strategies of the first player by $\{1, \dots, m\}$, and that of the second player by $\{m + 1, \dots, m + n\}$ (instead of $\{1, \dots, n\}$). This has the advantage that we can easily distinguish between strategies of the two players. When a strategy profile (x, y) is chosen, the payoff to player one is $x^\top Ay$, while that to player two is $x^\top By$. A strategy x of player one is a *best reply* to a strategy y of player two if it gives maximal expected payoff to player one, i.e. if we have that for all other strategies \bar{x} of the first player

$$x^\top Ay \geq \bar{x}^\top Ay$$

Player two's strategy y is a best reply to x if the analogous condition holds for the payoffs of player two. The *support* of a mixed strategy z is the set of all pure strategies played with positive probability; we denote it by $\text{supp}(z)$. The *best reply condition* (Nash, 1951) states that a mixed strategy x is a best reply to y if and only if every pure strategy in the support of x is a best reply to y . A strategy profile (x, y) is a Nash equilibrium if and only if x and y are best replies to each other.

A very useful geometric tool in the study of bimatrix games are the best reply polyhedron and best reply polytope. A *polyhedron* is a subset of \mathbb{R}^d that is a finite intersection of closed half-spaces; a *polytope* is a bounded polyhedron. Equivalently, a polytope is the convex hull of a finite set of points in \mathbb{R}^d . An inequality $c^\top x \leq \alpha$ is *valid* for a polyhedron P if it holds for all points x in P . A set $F \subset P$ is a *face* of P if there is a valid inequality $c^\top x \leq \alpha$ such that $F = P \cap \{x \mid c^\top x = \alpha\}$. The *dimension* of a face F of P is the dimension of its affine hull; F has dimension d if and only if F contains $d + 1$, but no more, affinely independent points. The 0- and 1-dimensional faces are called *vertices* and *edges*, respectively. For a polyhedron of dimension d , the faces of dimension $d - 1$ are called *facets*. A d -polytope is a d -dimensional polytope. The *edge graph* of a polytope consists of the vertices of the polytope, connected by its edges. To every polytope P we can assign a partially ordered set, its *face lattice*. It consists of the faces of P , partially ordered by inclusion. Two polytopes are called *combinatorially equivalent* if there is a bijection between their faces in each dimension that preserves face incidences, or in other words, if their face lattices are isomorphic. Two polytopes are called *affinely (linearly) isomorphic* if there is an affine (linear) map that induces a bijection between the polytopes. Two polytopes which are affinely isomorphic are, in particular, combinatorially equivalent (while the converse is generally wrong).

Given a bimatrix game (A, B) , the *best reply polyhedra* for player one and two are defined as

$$H_1 = \{(x, v) \in \mathbb{R}^m \times \mathbb{R} \mid B^\top x \leq v\mathbf{1}, x \geq \mathbf{0}, \mathbf{1}^\top x = 1\}$$

$$H_2 = \{(y, u) \in \mathbb{R}^n \times \mathbb{R} \mid Ay \leq u\mathbf{1}, y \geq \mathbf{0}, \mathbf{1}^\top y = 1\}$$

These polyhedra are the upper envelopes of the best reply function, which assigns to each strategy of a player the other player's payoff from his best reply (the best reply might not be unique, but the best reply payoff is). It is useful to *label* the

inequalities that define H_1 , as follows: For i in $\{1, \dots, m\}$, the inequality $x_i \geq 0$ has label i , while for j in $\{m+1, \dots, m+n\}$, the inequality $e_j^\top B^\top x \leq v$ has label j . This induces a labelling on the relative boundary of H_1 , where a point carries the labels of all inequalities that are binding in that point. The points on the relative boundary of H_2 are labelled in an analogous way.

We can use the projection p from H_1 onto player one's strategy space Δ_m to transfer this labelling to the strategy space. A point $x \in \Delta_m$ gets all labels that occur at some point in $p^{-1}(x)$. We get a subdivision of Δ_m into several labelled regions, where a point x has as labels the pure strategies in $\{1, \dots, m\}$ that are unplayed at x , and the pure strategies in $\{m+1, \dots, m+n\}$ that are the other player's best replies to x . We call the set of strategies of player one that have a certain label $j \in \{m+1, \dots, m+n\}$ of player two the *best reply region* with label j . Player two's strategy space Δ_n is subdivided and labelled in an analogous way.

It is obvious from the best reply condition that a strategy pair is a Nash equilibrium if and only if every strategy by either player which is not a best reply to the other player's strategy remains unplayed. Hence a strategy pair (x, y) is a Nash equilibrium if and only if it is *completely labelled* as a point in $\Delta_m \times \Delta_n$, i.e. if every label in $\{1, \dots, m, m+1, \dots, m+n\}$ occurs as a label of x or y . In this way the labelling of the strategy spaces of the two players can be used to visualize Nash equilibria of low dimensional bimatrix games.

The polyhedra H_1 and H_2 are unbounded, which makes them difficult to handle. However, they can be converted into polytopes, the *best reply polytopes*, which essentially are combinatorially equivalent to H_1 and H_2 (or more precisely, projectively equivalent; for a definition of projective equivalence see Chapter 4.8). Consider the polyhedra

$$P_1 = \{\bar{x} \in \mathbb{R}^m \mid \bar{x} \geq 0, B^\top \bar{x} \leq \mathbf{1}\}$$

$$P_2 = \{\bar{y} \in \mathbb{R}^n \mid \bar{y} \geq 0, A\bar{y} \leq \mathbf{1}\}$$

If these polyhedra are bounded, they are polytopes, called *best reply polytopes*. To achieve boundedness, it suffices to assume that the entries in the payoff matrices A and B are positive. This can be done without loss of generality since a constant can be added to the payoffs of any strategic-form game without changing the structure of the game. H_1 is in bijection with $P_1 \setminus \mathbf{0}$, via the map $(x, v) \mapsto x/v$. This map

is nonlinear but preserves binding inequalities, and therefore the face incidences. Hence from the labelling of H_1 , we obtain a labelling of the binding inequalities of P_1 and then, as above, a labelling of the points on the relative boundary of P_1 . Similarly, we label the relative boundary of P_2 . Every completely labelled point of $P_1 \times P_2$ corresponds to a Nash equilibrium, except for the vertex $\mathbf{0}$, which carries all labels because all strategies are unplayed, but does not correspond to any strategy profile. We call this vertex the *artificial equilibrium*.

A game is called *nondegenerate* if no point in Δ_m has more than m labels, and no point in Δ_n has more than n labels (see von Stengel (2002) for equivalent definitions of nondegeneracy). Equivalently, no strategy z of either player can have more than $|\text{supp}(z)|$ pure best replies. For a nondegenerate game, the polytopes P_1 and P_2 are *simple*, i.e. every vertex of P_1 is contained in exactly m facets, and similarly for P_2 . This implies that in a nondegenerate game, every vertex of P_1 has exactly m labels. Two adjacent vertices of P_1 share $m - 1$ labels, namely the labels of the edge connecting them. A similar observation holds for P_2 .

The best-known algorithm for the computation of a Nash equilibrium in a nondegenerate bimatrix game is the *Lemke-Howson algorithm* (Lemke and Howson, 1964). It is a complementary pivoting method that walks along the edges of the best reply polytopes. Denote by G the edge graph of the product polytope $P_1 \times P_2$. The vertices of G are given by pairs of vertices of P_1 and P_2 , while the edges of G are given by pairs of a vertex of P_1 and an edge of P_2 , or an edge of P_1 and a vertex of P_2 . The labellings of P_1 and P_2 induce a labelling of the vertices of G . Choose a label $k \in \{1, \dots, m + n\}$. A vertex of G is called *k-almost completely labelled* if it has every label apart from possibly k .

If we delete all vertices from the graph that are not k -almost completely labelled (and all corresponding edges) the new graph G_k contains all completely labelled vertices, which correspond to Nash equilibria, and all vertices that have as missing label k . Since the game is nondegenerate, adjacent vertices in G always have $m + n - 1$ labels in common. Hence the completely labelled vertices of G_k have degree 1 (which means that they have only one adjacent vertex in the graph), while the vertices with missing label k have degree 2 (they have two adjacent vertices). This implies that G_k consists of paths and cycles, where the endpoints of paths are completely labelled vertices. Starting at the vertex corresponding to the artificial

equilibrium, one can walk along the corresponding path in G_k , and at the end find a new completely labelled vertex, which must yield an equilibrium of the game.

For visualization, it is convenient to interpret the Lemke-Howson algorithm in terms of “picking up” and “dropping” labels in the two players’ strategy spaces. Instead of walking along edges of the polytope $P_1 \times P_2$, one can view the algorithm as tracing edges of the polyhedron $H_1 \times H_2$, with “infinity” as starting point. One can then project the Lemke-Howson path from $H_1 \times H_2$ onto the strategy space. The terminology of vertices and edges can be taken straightforwardly to the strategy space: A vertex in the strategy space Δ_m of player one is a point with m given labels, and an edge is a nonempty set of points with $m - 1$ given labels. Vertices and edges in the other player’s strategy space Δ_n are defined analogously. By definition, the (projected) Lemke-Howson algorithm moves alternately in the two players’ strategy spaces, jumping from vertex to vertex along an adjoining edge.

To be more precise, assume that the missing label k belongs to player one. The projected Lemke-Howson path with missing label k starts at the vertex (e_k, e_j) , where j is the best reply to k (which by nondegeneracy is unique). The vertex e_j has all labels in $\{m + 1, \dots, m + n\}$, apart from j , and one extra label i in $\{1, \dots, m\}$, which is player one’s best reply to j . If $i = k$, then (e_k, e_j) is an equilibrium. If $i \neq k$, the label i must have been present at player one’s vertex e_k (because the only missing label along the paths is k), hence can be dropped by player one in the next step. Player one then walks along the edge given by all labels present at e_k , apart from label i , until he reaches a new vertex, where he picks up a new label, which player two can then drop. This way, the players take turns in picking up and dropping labels, until an equilibrium is reached. As an example, consider the Lemke-Howson path with missing label 3 for the following game:

$$A = \begin{pmatrix} 4 & 4 & 4 \\ 0 & 0 & 6 \\ 5 & 0 & 0 \end{pmatrix}, B = \begin{pmatrix} 6 & 12 & 0 \\ 0 & 4 & 0 \\ 8 & 0 & 13 \end{pmatrix} \quad (1.1)$$

In this example, the Lemke-Howson algorithm finds the equilibrium $(5/11, 0, 6/11)$, $(4/5, 0, 1/5)$. We have illustrated the path that the algorithm traces in the strategy space in Figure 1.1. We will come back to this example in Chapter 3.




Figure 1.1: The Lemke-Howson path with missing label 3 for example (1.1). The left simplex is player one's, the right one player two's. Player one's strategies are labelled 1-3, player two's have labels 4-6. The labels in the simplex mark the players' best reply regions. The labels outside mark the edges of the simplex where the corresponding strategy is unplayed. The square dot is the equilibrium that is found by the Lemke-Howson algorithm. The black arrows give the path of the algorithm, and are numbered in the order in which they occur.

We have now collected all the prerequisites that we require for the following chapters. Whatever else we need will be explained as we go along.

2

The index of an equilibrium component

2.1 Introduction

In this chapter, we give an exposition to the index, and present an algorithm for the computation of the index of an equilibrium component in degenerate games. The index is a topological notion which assigns to each connected component of Nash equilibria an integer which can be interpreted as an “orientation” of the equilibrium component. As explained in Chapter 1, the index is useful in a variety of contexts, particularly in the theory of equilibrium selection and refinement.

In nondegenerate games, where all equilibria are isolated, the index of an equilibrium can be easily computed, essentially as the sign of a suitable determinant (see Definition 4.3). However, for equilibrium components in degenerate games, there is no such explicit formula. A general method to calculate the index of an equilibrium component works as follows: Choose a nondegenerate perturbation of the game, compute the equilibria of the perturbed game, and add up the indices of those equilibria that are close to the given component (Demichelis and Germano, 2000; Ritzberger, 2002). However, this approach leads to several complications: First, it is necessary to actually perturb the game. Second, we need to decide on when an equilibrium of the perturbed game is close enough to the original component to warrant being included in the calculation of its index. Due to these drawbacks,

research has focused on games with a simple structure, like outside option games (see, for example, Hauk and Hurkens, 2002, von Schemde, 2005), where the index of a given component is easily computable using the fact that in any bimatrix game, the indices of all equilibrium components have to add up to $+1$.

However, for most equilibrium components, such a line of argument is not sufficient, hence a more general method is needed. The algorithm that we suggest is, in a sense, a simplification of the perturbation method described above. However, since our algorithm does not require any explicit perturbations, it avoids the disadvantages of this method.

Our algorithm works as follows: For a given equilibrium component in a degenerate game we consider its extreme equilibria, and assign to each of these a new integer, which we call its *lex-index*. The crucial ingredient for our algorithm is the following, intuitively appealing result:

Theorem 2.1. *The index of an equilibrium component is the sum of the lex-indices over all extreme equilibria of this component.*

Since the lex-index of an extreme equilibrium is easily computable, this result gives immediately rise to a “perturbation-free” algorithm for the computation of the index of a component of Nash equilibria. For its proof, we use the standard procedure for index computation, as described above: We perturb the game and add up the indices of equilibria near a given component. For our purpose, we choose a lexicographic perturbation. Under this perturbation, every extreme equilibrium decomposes into a finite number of isolated equilibria of the perturbed game, or vanishes. The lex-index of an extreme equilibrium is defined as the sum of the indices of all equilibria of the perturbed game that originate from this particular extreme equilibrium. Using this definition of the lex-index, Theorem 2.1 follows immediately.

The advantage of our method is that the lex-index can be calculated using a very simple approach. Both the equilibria of the lexicographically perturbed game and their indices can be easily computed. Also, for every equilibrium of the perturbed game it is immediately clear which extreme equilibrium of the original game it comes from. More precisely, the equilibria of the original game solve a linear complementarity problem (LCP), i.e. an optimization problem with linear constraints

and a complementarity condition (for precise definitions, see Section 2.3). Every extreme equilibrium of the original game has one or multiple representations by “bases” of this LCP. Every basis that is both “lexico-feasible” and “complementary” corresponds to a unique equilibrium of the perturbed game (the first property, lexico-feasibility, is needed in order to maintain feasibility of the basis, the second, complementarity, to keep the equilibrium property in the perturbed game). The index of this equilibrium can be computed from the payoff matrices using a suitable lexico-rule. This in turn implies that the lex-index of an extreme equilibrium of the original game can be easily calculated, by adding up the indices of those equilibria of the perturbed game that come from the bases which represent this particular extreme equilibrium. It also means that we can avoid explicitly computing the equilibria of the perturbed game, and their indices. In this sense, the concept of lex-index, and with it our description of the index of a component in Theorem 2.1, does not rely on perturbations of the game.

The structure of the chapter is as follows: In Section 2.2, we give a short exposition to various definitions and properties of the index, and summarize the theoretical foundations of our algorithm. In Section 2.3, we describe how to perturb a game lexicographically, and analyze the equilibrium structure of the perturbed game. We use these perturbations to develop the concept of lex-index, and prove Theorem 2.1 in Section 2.4. We also reformulate this theorem as an explicit algorithm, see Algorithm 2.12.

2.2 The index

Let (A, B) be an $m \times n$ bimatrix game. Since adding a positive constant to the matrices does not change the structure of the game, we can assume without loss of generality that $A, B > 0$ for the remainder of this chapter. Recall that the support $\text{supp}(z)$ of a strategy z of either player is the set of pure strategies played with positive probability, and $|\text{supp}(z)|$ the number of strategies in the support of z . The game is called nondegenerate if every strategy x of the first player has at most $|\text{supp}(x)|$ pure best replies, and similarly for the second player. For an equilibrium (x, y) , denote by A_{xy} and B_{xy} the matrices obtained from A and B by deleting all rows that are not contained in the support of x , and all columns that are not con-

tained in the support of y . In a nondegenerate game, it is straightforward that in an equilibrium, both players use strategies of equal support size, hence these matrices must be square. Moreover, they must be nonsingular, since any linear dependency between rows or columns of the matrices A_{xy} or B_{xy} can be used to reduce the support of one of the strategies without changing the set of its pure best replies, contradicting nondegeneracy (von Stengel, 2002). The *index* of an equilibrium in a nondegenerate game can be defined as follows (Shapley, 1974):

Definition 2.2. *Let (x, y) be a Nash equilibrium of a nondegenerate bimatrix game (A, B) , where $A, B > 0$. The index of (x, y) is defined as*

$$(-1)^{|\text{supp}(x)|+1} \text{sign det}(A_{xy}) \text{det}(B_{xy}). \quad (2.1)$$

This definition can be extended to games with potentially non-positive entries: Just add a sufficiently large constant to the game and then define the index using equation (2.1). This is well-defined, due to part (b) of the next Proposition.

Actually, in Shapley's definition of the index, the sign of the index is reversed. However, the above sign convention has been shown to be more convenient. Shapley's main result, which motivated his definition of the index, holds regardless of the chosen sign convention: Equilibria at opposite ends of a Lemke-Howson path have opposite index.

The following proposition (von Schemde and von Stengel, 2008, Proposition 2) collects some well-known properties of the index.

Proposition 2.3. *In a nondegenerate game, the index of a Nash equilibrium*

- (a) *is $+1$ or -1 ;*
- (b) *does not change when adding a positive constant to all payoffs;*
- (c) *only depends on the payoffs in the support of the equilibrium;*
- (d) *does not depend on the order of the players' pure strategies;*
- (e) *is $+1$ for any pure-strategy equilibrium;*
- (f) *the sum of the indices over all equilibria is $+1$.*

Most of these properties are obvious from the definition, others require some more work, see von Schemde and von Stengel (2008).

In degenerate games, the definition of the index is more involved since Definition 4.3 can no longer be applied. There are several ways to extend the definition of the index to degenerate games and components of equilibria, most of which rely on the notions of global and local degree of a continuous map. The degree is a quite advanced topological tool, which is based on the concept of orientations of manifolds. We give a short exposition to the local and global degree of a continuous map using standard results and methods from algebraic topology (see, for example, Dold, 1980, Sections IV.5 and VIII.4, or Hatcher, 2001). An introduction to the intuition behind the concept of degree can be found in Demichelis and Germano (2000).

To be mathematically precise, we will have to use homology groups (with integer coefficients). We try, however, to give some geometric interpretation along the way that should be accessible to a reader without any knowledge in algebraic topology. The rough idea behind the degree is the following: Consider a topological d -manifold X (i.e. a topological space that is locally homeomorphic to \mathbb{R}^d). Intuitively, the manifold is orientable if we can choose “local” orientations around every point x in X that are compatible globally, i.e. “glue together nicely”. Hence a (global) orientation is a collection of local orientations that “fit together”. Now the *global degree* of a continuous map f between compact connected orientable d -manifolds X and Y measures what happens to the global orientation when we apply f . The *local degree* of f around a point x measures what happens to the local orientation around x . For our purposes, it is not necessary to understand the precise definition of the index as a local degree. The reader preferring to avoid the topological bit can skip the following paragraphs and accept Proposition 2.4 as a definition of the index of an equilibrium component. From Proposition 2.4 onwards we will not need the notion of local or global degree for the remainder of this thesis.

We formalize the intuitive definition of the degree in terms of homology groups. We assume the reader to be familiar with the concept of homology; otherwise, the following can be taken as intuitive (albeit imprecise) “definitions”: For a topological space X and $i \in \mathbb{N}$, the *i th homology group* $H_i X$ essentially measures the i -dimensional shape of X . For a subset W of X , the *i th relative homology group* $H_i(X, W)$ encodes the relationship between the homology groups $H_i W$ and $H_i X$, in a sense made precise by the “long exact homology sequence”. The intuitive

idea behind the relative homology is to ignore everything that is contained in the subspace W . Every continuous map $f : X \rightarrow Y$ between topological spaces induces group homomorphisms $H_i f : H_i X \rightarrow H_i Y$. If $W \subset X$ and $Z \subset Y$ such that $f(W) \subset Z$, the map f also induces homomorphisms $H_i f : H_i(X, W) \rightarrow H_i(Y, Z)$. A central observation is that for a d -manifold X and a point x in X , the relative homology group $H_d(X, X \setminus \{x\})$ is isomorphic to the group \mathbb{Z} of integers. The intuition behind this is as follows: Relative homology is “local” (in a sense made precise by “excision”, a very useful tool in homology theory). Since a d -manifold X looks locally like \mathbb{R}^d , the group $H_d(X, X \setminus \{x\})$ is isomorphic to $H_d(\mathbb{R}^d, \mathbb{R}^d \setminus \{0\})$, which is a free group generated by a d -cycle (i.e. a d -simplex or “topological ball”) around 0.

Hence a generator of $H_d(X, X \setminus \{x\})$ can be interpreted as the choice of a cycle around the point x . Intuitively, such a cycle, understood as a d -simplex around x , orients X locally around x , by giving the space “direction”. Hence it makes sense to define a *local orientation* of X at a point x to be a choice of generator of $H_d(X, X \setminus \{x\})$. Equivalently, a local orientation around x is the choice of an isomorphism $H_d(X, X \setminus \{x\}) \xrightarrow{\sim} \mathbb{Z}$. A *global orientation* is a function that assigns to each point x a local orientation o_x , in a locally consistent way. A manifold is *orientable* if it has at least one global orientation.

Orientability of a compact, connected d -manifold X implies that the top homology group $H_d X$ is isomorphic to \mathbb{Z} . A global orientation of X then corresponds to the choice of such an isomorphism, or equivalently a choice of a generator of $H_d X$. This generator is usually called the *fundamental cycle* of X . A continuous map $f : X \rightarrow Y$ between compact, connected, oriented d -manifolds induces a group homomorphism $H_d f$ on the top cohomology groups, which must correspond to multiplication by an integer. This integer is called the (*global*) *degree* of f . It counts the “multiplicity” with which $H_d f$ maps the fundamental cycle on X to the fundamental cycle on Y , hence quantifies how the global orientation changes under f .

Similarly to a global degree, a continuous map $f : X \rightarrow Y$ between oriented manifolds induces a local degree in the following way: Assume we have a point y in the range of f whose preimage $f^{-1}(y)$ consists of only one point x . Then f induces a homomorphism $H_d f : H_d(X, X \setminus \{x\}) \rightarrow H_d(Y, Y \setminus \{y\})$. As seen above, both of these homology groups are isomorphic to \mathbb{Z} , where the isomorphism depends on

the chosen orientation. So $H_d f$ corresponds to multiplication by an integer, which is called the *local degree* of f over y . This degree measures the number of cycles around y obtained from a cycle around x under $H_d f$. If we interpret the cycle as a local orientation around x , the local degree quantifies how this local orientation changes under f . This interpretation coincides with an equivalent definition of the local degree for differentiable maps of differential manifolds: In the situation above, assume that additionally, f is differentiable, and x a regular value of f . Then the local degree of f over y is the sign of the determinant of the Jacobian at x . In this sense, the local degree indicates if the map is locally orientation-preserving or -reversing.

The local degree can be extended to the case where $f^{-1}(y)$ is compact, following a similar concept. A global orientation on X induces a local orientation along every compact subset of X . On the d th relative homology groups, this orientation along the compact set is mapped to an integer multiple of the local orientation around y . This integer is the local degree of f over y . An important property that we need later is “locality” of the local degree: The local degree of f over y does not change if we restrict f to some neighborhood of $f^{-1}(y)$, since relative homology groups are “local”.

Another useful property of the local degree is additivity. Assume X is a finite union of open sets X_i such that the sets $f^{-1}(y) \cap X_i$ are pairwise disjoint. Then the local degree of f over y is the sum of the local degrees of $f|_{X_i}$ over y . This means that the local degree of f over y is composed of the local degrees of “localized” (i.e. restricted) versions of f .

A crucial notion in the context of degree is that of a proper map. A continuous map is *proper* if the inverse image of every compact set is again compact. If X and Y are oriented manifolds, and Y is connected, then for a proper map $f : X \rightarrow Y$ the local degree over y is the same for every y in Y . (Moreover, if the manifolds are both compact and connected, this degree equals the global degree of f as defined earlier.)

The degree of a continuous map becomes useful in game theory because the index of an equilibrium component can be defined in terms of degree. There are several ways of doing this. For example, the equilibria of a bimatrix game can be represented as rest points, i.e. zeros, of certain vector fields, called *Nash fields*.

The index of an equilibrium component can then be defined in dynamic terms as the “Poincaré index” of the corresponding component of zeroes of such a Nash field (see Ritzberger, 1994, or Demichelis and Germano, 2000). The index of a component of rest points of a Nash field f is the global degree of the map $f/||f||$, restricted to a small “sphere” around the component of rest points. It can also be seen as the local degree of f , restricted to a small neighborhood of the component, over 0. For regular equilibria, it is the sign of the determinant of the corresponding Jacobian. In this interpretation, the index of a rest point carries information about its dynamic stability. Equivalently, if we have a *Nash map* f , i.e. a continuous map on the strategy space whose fixed points are the equilibria of the game, then the index of the equilibrium component is the local degree of $id - f$, restricted to a neighborhood of the component, over zero (by id we denote the identity map). This definition does not depend on the choice of Nash map, and for nondegenerate games yields indeed Shapley’s definition (Govindan and Wilson, 1997b).

If understood as a fixed point index, the index of an equilibrium component carries information about the “stability” or “essentiality” of a component. A component of fixed points of a continuous map is called essential if the component does not vanish under continuous perturbations of the underlying map. This is equivalent to its index being nonzero (O’Neill, 1953). Hence an equilibrium component of nonzero index will be stable in the sense that whatever Nash map we choose, every perturbation of that map will have a fixed point close to that equilibrium.

In game theoretic terms, this “essentiality” of an equilibrium component with nonzero index translates into a version of hyperstability, a concept that goes back to Kohlberg and Mertens (1986). Recall that a pure strategy of one player is called *redundant* if that player has an *equivalent strategy*, i.e. a convex combination of his other pure strategies that gives the same expected payoff against any strategy of the other player. From a game (A, B) we get the *reduction* of the game by deleting all redundant strategies. Two games are called *equivalent* if they yield the same reduced game. Govindan and Wilson (2005) call an equilibrium component of a game (A, B) *hyperstable* if for every equivalent game (A^*, B^*) and every neighborhood U of the component, there is a $\delta > 0$ such that every δ -perturbation of (A^*, B^*) has an equilibrium that is equivalent to some strategy in U . They call an equilibrium *uniformly hyperstable* if the hyperstability condition is met uniformly,

i.e. if δ can be chosen independently of the equivalent game. Govindan and Wilson (2005) proved that an equilibrium component has nonzero index if and only if it is uniformly hyperstable, answering a long-standing open question as to what the suitable game-theoretic counterpart of topological essentiality should be. For a detailed discussion of the link between essentiality and the theory of equilibrium selection see Govindan and Wilson (1997b).

For computational purposes, we will need yet another definition of the index as the local degree of the projection from the graph of the equilibrium correspondence to the space of games. More precisely, consider $\mathbb{R}^{m \times n} \times \mathbb{R}^{m \times n}$ as the space of $m \times n$ bimatrix games. Let \mathcal{E} be the graph of the equilibrium correspondence over this space of games, i.e. the correspondence which maps each game to its set of equilibria. Consider the projection map $p : \mathcal{E} \rightarrow \mathbb{R}^{m \times n} \times \mathbb{R}^{m \times n}$; by Kohlberg and Mertens (1986), it is homotopic to a homeomorphism. This means that we can orient E such that p has global degree 1.

Consider an equilibrium component C of a bimatrix game (A, B) , and let U be a neighborhood of C in \mathcal{E} that “separates” C from all other equilibria of (A, B) . By this we mean that U does not contain any other equilibria of (A, B) apart from C . Then the index of C coincides with the local degree of the projection map p , restricted to U , over the game (A, B) (see Govindan and Wilson, 1997a, for bimatrix games, and Demichelis and Germano, 2000, for general strategic-form games). Note that by “locality” of the local degree, this definition does not depend on the choice of U . From the additivity of the local degree it follows easily that for non-degenerate as well as for degenerate games, the sum of the indices over all equilibrium components equals +1. However, the index of an equilibrium component in degenerate games is no longer restricted to $\{+1, -1\}$; any integer can occur as index (Govindan et al., 2003).

It is a well known fact that in order to compute the index of an equilibrium component, we can perturb the game slightly, and add up the indices of the equilibria nearby, see, for example, Demichelis and Germano (2000) or Ritzberger (2002). However, since this “perturbation-method” is the foundation of our algorithm, we would like to give and prove a more precise statement, which we suspect to be well-known but for which we have not found a reference:

Proposition 2.4. *Suppose we are given an equilibrium component C of an $m \times n$*

bimatrix game (A, B) . For any neighborhood U of C in the graph of the equilibrium correspondence \mathcal{E} whose closure does not contain any other equilibria of (A, B) , there is a neighborhood V of (A, B) in the space of $m \times n$ -games such that the following holds: For every game $(A', B') \in V$, the sum of the indices of the equilibrium components of (A', B') that are contained in U equals the index of C .

Proof. The proposition is essentially due to the fact that the projection is locally proper, which implies that the local degree is “locally independent” of the game chosen. Coupled with additivity of the local degree, the proposition follows. In more detail, the proof works as follows: Consider the restriction of the projection $p|_U : U \rightarrow \mathbb{R}^{m \times n} \times \mathbb{R}^{m \times n}$. The index of C is the local degree of $p|_U$. The local degree of a proper map is constant (if the range is connected) (Dold, 1980, IV.5.12 or VIII.4.5). Hence we just need to find a connected open neighborhood V of the game (A, B) such that $p|_U$ is “proper over V ”, i.e. such that $p|_{U \cap p^{-1}(V)}$ is proper.

Let $\bar{U}_\delta(C)$ be the compact δ -neighborhood of C in the equilibrium graph, and choose $\delta > 0$ small enough such that $\bar{U}_\delta(C)$ is contained in U . Choose V such that for some $\delta > 0$, $(p|_U)^{-1}(\bar{V})$ is contained in $\bar{U}_\delta(C)$ (where \bar{V} is the closure of V). Such a choice of V is possible since the equilibrium correspondence is upper hemicontinuous, or more precisely, since its graph is closed (and because \bar{U} does not contain any equilibria apart from C). For this choice of V , let $U' = U \cap p^{-1}(V)$. By construction of V , $p|_{U'} : U' \rightarrow V$ is proper. This, together with localness of the degree, implies that the local degree of $p|_U$ is “constant over V ”, by which we mean that it is the same over every (A', B') in V . By additivity (Dold, 1980, Theorem IV.5.8 or Proposition VIII.4.7), the latter degree is just the sum of the indices of the equilibria of (A', B') that are contained in U . \square

The version of the index in Proposition 2.4 is most useful for computations, since it allows for a computation of the index of an equilibrium component based on perturbations of the game. However, as explained in the introduction, direct application of this perturbation-approach leads to several complications. In order to arrive at a simpler, “perturbation-free” method, we use lexicographic perturbations as the base of our algorithm, which we develop in the next two sections.

2.3 A lexicographically perturbed game

Lexicographic perturbations are used to solve degenerate linear programs and linear complementarity problems by making them nondegenerate (Dantzig, 1963). In game theory, these perturbations have been used to resolve degeneracies in the Lemke-Howson algorithm (Lemke and Howson, 1964), and the van den Elzen-Talman algorithm (von Stengel et al., 2002). We use a slight variation of the concept of lexicographic perturbations to turn a degenerate game into a nondegenerate one. In the following section, we use this concept of lexicographic perturbation of a degenerate game to compute the indices of its equilibrium components.

Recall that the set of equilibria of a bimatrix game can be understood as the solutions to a linear complementarity problem (which is a special case of a quadratic programming problem with linear constraints and a complementarity condition). Given a bimatrix game (A, B) , the equilibria (x, y) of the game are in one-to-one correspondence with the solutions to the following set of equations and inequalities (von Stengel, 1996):

$$\begin{aligned} \mathbf{1}^\top y &= 1 \\ -u\mathbf{1} + Ay + r &= 0 \\ y, r &\geq 0 \end{aligned} \tag{2.2}$$

and

$$\begin{aligned} \mathbf{1}^\top x &= 1 \\ -v\mathbf{1} + B^\top x + s &= 0 \\ x, s &\geq 0 \end{aligned} \tag{2.3}$$

such that

$$x^\top r = 0 = y^\top s \tag{2.4}$$

The variables u and v give the payoffs to the respective players, and r and s are slack variables that measure how far from being optimal a strategy is.

The systems (2.2) and (2.3) are linear functions of some constrained and unconstrained variables, i.e. they are of the form

$$D(z, z') = b, z' \geq 0$$

for some matrix D , a set of *unconstrained* variables z and a set of *constrained* variables z' . In our case, the only unconstrained variables are u and v , respectively; all other variables are constrained to be nonnegative. D is a $k \times k'$ matrix, where $k' > k$, and has maximal possible rank k : In the first case, D is of the form

$$\begin{pmatrix} 0 & \mathbf{1}^\top & 0 \\ -\mathbf{1} & A & I \end{pmatrix} \quad (2.5)$$

and in the second

$$\begin{pmatrix} 0 & \mathbf{1}^\top & 0 \\ -\mathbf{1} & B^\top & I \end{pmatrix} \quad (2.6)$$

where I denotes the identity matrix of suitable size. We will need the following standard terminology of linear programming and its extensions (see, for example, Dantzig, 1963): A solution to the system $D(z, z') = b$ is called *feasible* if it satisfies the nonnegativity constraints $z' \geq 0$. A feasible solution is called *extreme* if it cannot be written as a proper convex combination of two other feasible solutions. A solution to $D(z, z') = b$ is called *basic* if the columns of D that correspond to the unconstrained and the nonvanishing constrained variables, are linearly independent. A *basis* or *basic set of variables* consists of any set of k variables, containing all unconstrained variables, such that the square matrix given by the columns of D corresponding to these variables is non-singular, i.e. those columns form a basis of \mathbb{R}^k . To every such basis we can assign a basic solution of the equation $D(z, z') = b$ by setting the non-basic variables to zero and solving for the basic variables (the solution is then unique). By abuse of terminology we call a basis feasible if the corresponding basic solution is feasible. The system $D(z, z') = b, z' \geq 0$ is called *non-degenerate* if in every basic feasible solution, all constrained basic variables have positive value. By von Stengel (1996), degeneracy of a bimatrix game (given by positive matrices) is equivalent to degeneracy of the corresponding systems (2.2) and (2.3).

A central role in our algorithm is played by *extreme equilibria*. A Nash equilibrium is called extreme if it cannot be written as a proper convex combination of other Nash equilibria of the game, i.e. if it gives rise to extreme solutions of (2.2) and (2.3). Extreme equilibria have the following property:

Proposition 2.5. *An equilibrium (x, y) of a bimatrix game (A, B) is extreme if and only if the corresponding solutions of the systems (2.2) and (2.3) are basic.*

Proof. This follows directly from the following standard result: A point (z, z') in $Z = \{(z, z') \mid D(z, z') = b, z' \geq 0\}$ is an extreme point of Z if and only if (z, z') is a basic feasible solution. For systems where all variables are constrained, a proof of this result is given in Dantzig (1963, Theorem 7.1.3); the same proof also works for systems with unconstrained variables. \square

We now have all the terminology that we need to describe lexicographic perturbations of bimatrix games. For an $m \times n$ bimatrix game (A, B) and $\varepsilon > 0$, define a perturbed game $(A(\varepsilon), B(\varepsilon))$ where $A(\varepsilon) = A - E(m, n)$, $B(\varepsilon) = B - E(n, m)^\top$, where $E(m, n)$ is the $m \times n$ matrix given by

$$E(m, n) = \underbrace{\begin{pmatrix} \varepsilon & \dots & \varepsilon \\ & \dots & \\ \varepsilon^m & \dots & \varepsilon^m \end{pmatrix}}_{n \text{ columns}}$$

Perturbing the game in this way essentially corresponds to lexicographic perturbations of the corresponding systems (2.2) and (2.3): If we replace A by $A(\varepsilon)$ in (2.2), we get the following system of equations and inequalities:

$$\begin{aligned} \mathbf{1}^\top y &= 1 \\ -u\mathbf{1} + Ay + r &= (\varepsilon, \dots, \varepsilon^m)^\top \\ y, r &\geq 0 \end{aligned} \tag{2.7}$$

and similarly for system (2.3)

$$\begin{aligned} \mathbf{1}^\top x &= 1 \\ -v\mathbf{1} + B^\top x + s &= (\varepsilon, \dots, \varepsilon^n)^\top \\ x, s &\geq 0 \end{aligned} \tag{2.8}$$

A solution to those two perturbed systems yields an equilibrium of $(A(\varepsilon), B(\varepsilon))$ if and only if it satisfies the complementarity condition (2.4), i.e. if $x \perp r$ and $y \perp s$.

Since the solutions of (2.7) are the same as of the system (2.2) with A replaced by $A(\varepsilon)$, these systems also have the same sets of extreme feasible solutions. As seen in the proof of Proposition 2.5, the extreme feasible solutions correspond to feasible bases of the respective systems. Hence the feasible bases of the two systems

coincide. This may sound obvious, but is equivalent to the following seemingly non-trivial statement: A subset of $m + 1$ columns of

$$\begin{pmatrix} 0 & \mathbf{1}^\top & 0 \\ -\mathbf{1} & A & I \end{pmatrix} \quad (2.9)$$

is linearly independent and feasible (i.e. the corresponding basic solution is feasible) if and only if the corresponding subset of

$$\begin{pmatrix} 0 & \mathbf{1}^\top & 0 \\ -\mathbf{1} & A(\varepsilon) & I \end{pmatrix} \quad (2.10)$$

is linearly independent and feasible. Similarly, the basic feasible solutions of the system (2.8) will be the same as those to the system (2.3), with B replaced by $B(\varepsilon)$. We will use this property implicitly several times.

Note that the perturbations in (2.7) and (2.8) are not standard lexicographic perturbations of (2.2) and (2.3), in that the first row in both systems remains unperturbed. However, these perturbations still lead to nondegenerate systems. The intuition behind this is that the first line of both systems is never degenerate since the degeneracies of a game are contained in the matrices A and B , respectively.

Proposition 2.6. *For small enough $\varepsilon > 0$, the game $(A(\varepsilon), B(\varepsilon))$ is nondegenerate.*

It is well-known that a lexicographically perturbed system of equations is nondegenerate. We slightly adapt the standard proof of this fact to suit out non-standard perturbations. Recall that a nonzero vector is called *lexicopositive* if its first nonzero entry is positive, and *lexiconegative* otherwise.

Proof. According to von Stengel (1996), we need to prove that in any basic feasible solution to the system (2.7) all basic variables are positive (the corresponding statement for system (2.8) will follow by analogy). Now (2.7) is of the form $D(u, y, s) = (1, 0, \dots, 0)^\top + (0, \varepsilon, \dots, \varepsilon^m)^\top$ for the matrix D given in (2.5). For any basis of the column space of D , denote by S the submatrix of D given by the basic columns. Then the corresponding basic solution is given by $S^{-1}(1, 0, \dots, 0)^\top + S^{-1}(0, \varepsilon, \dots, \varepsilon^m)^\top$. If we denote the columns of S^{-1} by $[\bar{s}_1, \dots, \bar{s}_{m+1}]$, we get that the basic solution is given by $\bar{s}_1 + \varepsilon \bar{s}_2 + \dots + \varepsilon^m \bar{s}_{m+1}$. The matrix S^{-1} can have no zero row since it has full rank. If a row of S^{-1} is lexicopositive, the corresponding basic variable is positive for ε small enough. If the row is lexiconegative, the basic

variable is negative (hence the corresponding solution will be unfeasible). Hence in every basic feasible solution to (2.7), the basic variables will be positive. \square

In order to use the “perturbation method” from Proposition 2.4 to calculate the index of an equilibrium component of a degenerate game (A, B) , we need to understand the equilibrium structure of $(A(\varepsilon), B(\varepsilon))$. It turns out that for this, we need to gain some deeper understanding of bases of the systems (2.2), (2.3), and their relationship to extreme equilibria of (A, B) . In the light of the proof to Proposition 2.6 we would like to remind the reader of the following terminology: A basis of the systems (2.2) or (2.3) is called *lexicofeasible* if for the corresponding basic matrix S , its inverse matrix S^{-1} has only lexicopositive rows, i.e. rows in which the first non-zero entry is positive. Moreover, we call a pair of bases of (2.2) and (2.3) *complementary* if the following condition holds: For any $1 \leq j \leq n$ we have that at least one of the variables y_j and s_j is nonbasic, and similarly for any $1 \leq i \leq m$, that at least one of the variables x_i and r_i is nonbasic. In particular, this implies that the corresponding basic solutions satisfy the complementarity condition (2.4).

In a nondegenerate game, every equilibrium corresponds to a unique pair of bases of the systems (2.2) and (2.3). By nondegeneracy of the system, this pair of bases will have to be complementary and lexicofeasible (where by abuse of terminology we mean a pair of bases to be lexicofeasible if both bases are). For an extreme equilibrium (x, y) in a degenerate game, however, there may be several bases for the system (2.2) that yield y as solution, and similarly there may be several bases of (2.3) that yield x as solution. Hence we may get several pairs of bases of (2.2), (2.3) that correspond to the extreme equilibrium (x, y) . In general, only a few of such pairs of bases will be both complementary and lexicofeasible. The next Proposition tells us how to tell the equilibria of $(A(\varepsilon), B(\varepsilon))$ from the complementary and lexicofeasible pairs of bases of the systems (2.2) and (2.3).

Proposition 2.7.

- (i) For $\varepsilon > 0$ sufficiently small, the equilibria of $(A(\varepsilon), B(\varepsilon))$ are in one-to-one correspondence with the complementary and lexicofeasible pairs of bases of (2.2), (2.3).
- (ii) Fix such a pair of bases. As ε tends to zero, the corresponding equilibrium of $(A(\varepsilon), B(\varepsilon))$ converges to an extreme equilibrium of (A, B) . This limit

equilibrium is given by the corresponding basic solutions of the unperturbed systems (2.2), (2.3).

Proof. By the proof of Proposition 2.6, for small ε , lexicofeasibility of a basis of (2.2) is equivalent to feasibility of the corresponding solutions to the system (2.7), and similarly for a basis of the system (2.3). Complementarity of the pair of bases ensures that condition (2.4) is met. Moreover the established correspondence is one-to-one, since in a nondegenerate game, every pair of complementary bases gives rise to a different equilibrium. This proves the first part of the Proposition.

As to the second part, for any pair of lexicofeasible complementary bases, the basic feasible solutions of the perturbed systems (2.7), (2.8) converge to a solution of the original systems (2.2), (2.3). The latter solution must then be feasible as well, and also satisfies the complementarity condition (2.4). Hence it corresponds to an equilibrium of (A, B) , which must be extreme due to Proposition 2.5. \square

From now on, when we consider the game $(A(\varepsilon), B(\varepsilon))$, let $\varepsilon > 0$ be sufficiently small for Propositions 2.6 and 2.7(i) to hold. Proposition 2.7 links every extreme equilibrium of the unperturbed game (A, B) to a set of equilibria of $(A(\varepsilon), B(\varepsilon))$, via certain bases of the system (2.2), (2.3) that yield that extreme equilibrium as a solution. Note that an extreme equilibrium in a degenerate game (A, B) may give rise to several complementary pairs of bases, none (or many) of which may be lexicofeasible. Consider, for example, the degenerate game

$$A = \begin{pmatrix} 1 & 1 \\ 1 & 2 \end{pmatrix} = B^\top \quad (2.11)$$

which has two pure Nash equilibria: $(1, 0), (1, 0)$ and $(0, 1), (0, 1)$. The first of these extreme equilibria gives rise to several bases: There are three basic sets of variables for system (2.2) that contain y_1 : $\{u, y_1, y_2\}$, $\{u, y_1, r_1\}$ and $\{u, y_1, r_2\}$. Analogously, (2.3) has three corresponding sets of basic variables. Except for the second, these bases are not lexicofeasible. The bases of the two systems can be combined to nine different pairs, among which only one is lexicofeasible, namely the pair given by the basic variables $\{u, y_1, r_1\}$ and $\{v, x_1, s_1\}$. However, this basis is not complementary. Hence by Proposition 2.7, there will be no equilibrium of $(A(\varepsilon), B(\varepsilon))$ close to $(1, 0), (1, 0)$. This is in line with the fact that for any positive ε , the equilibrium $(1, 0), (1, 0)$ will vanish in $(A(\varepsilon), B(\varepsilon))$ since it is strongly dominated.

2.4 Computing the index of an equilibrium component: The lex-index

Consider an equilibrium component C of a degenerate game (A, B) . Proposition 2.7 tells us how to find all the equilibria of the lexicographically perturbed game $(A(\varepsilon), B(\varepsilon))$ that are close C . By Proposition 2.4, we need to calculate the indices of these equilibria, and add them up in order to get the index of the component. The first part boils down to calculating determinants. It turns out that we will not have to actually compute the determinants of submatrices of $A(\varepsilon)$ and $B(\varepsilon)$, as we would have to if we wanted to use Definition 4.3 for the index computation. Since we chose a lexicographic perturbation, we will see that it suffices to calculate determinants related to the matrices A and B . Moreover, it turns out that for an equilibrium of $(A(\varepsilon), B(\varepsilon))$, its index will depend only on the corresponding complementary pair of lexicofeasible bases, as we will prove in our next Proposition. First, however, we need to introduce the following notation:

Definition 2.8. For a square $k \times k$ matrix M , denote by $[M \mid_i \mathbf{1}]$ the matrix obtained from M by replacing the i th column by the vector $\mathbf{1}$. Define $\zeta(M)$, called the sign of M , to be $+1$ if the vector

$$[\det(M), -\det([M \mid_1 \mathbf{1}]), \dots, -\det([M \mid_k \mathbf{1}))]$$

is lexicopositive, -1 if it is lexiconegative, and 0 if it vanishes.

The following proposition expresses the index of an equilibrium of $(A(\varepsilon), B(\varepsilon))$ in terms of signs of submatrices of A and B .

Proposition 2.9. For any lexicofeasible and complementary pair (α, β) of bases to (2.2), (2.3), define $|\beta|$ to be the number of variables x_i that are basic (which equals $|\alpha|$, the number of variables y_j that are basic, since the bases are complementary). Let $A_{\alpha\beta}$ and $B_{\alpha\beta}$ be the matrices obtained from A and B , respectively, by deleting all rows and columns corresponding to non-basic variables x_i and y_j .

For ε sufficiently small, (α, β) corresponds to an equilibrium of the perturbed game $(A(\varepsilon), B(\varepsilon))$. The index of this equilibrium is

$$(-1)^{|\beta|+1} \zeta(A_{\alpha\beta}^\top) \zeta(B_{\alpha\beta})$$

Before we prove Proposition 2.9, a useful observation:

Lemma 2.10. *For a $k \times k$ -matrix M and any k -vector ξ , the determinant of $M + [\xi, \dots, \xi]$ is*

$$\det(M) + \sum_{i=1}^k \det([M \mid_i \xi])$$

where $[M \mid_i \xi]$ denotes the matrix obtained from M by replacing the i th column with ξ .

Proof. The determinant is multilinear and vanishes if two columns are linearly dependent. □

Proof of Proposition 2.9. Since the game $(A(\varepsilon), B(\varepsilon))$ is nondegenerate, the support size of the equilibrium (x, y) is $|\beta|$. By Definition 4.3, the index of the equilibrium is the sign of

$$(-1)^{|\beta|+1} \det(A(\varepsilon)_{\alpha\beta}) \det(B(\varepsilon)_{\alpha\beta})$$

To compute the first determinant, abbreviate $A_{\alpha\beta}$ by \mathbf{A} , denote the support of x by $\{i_1, \dots, i_k\}$, where $k = |\beta|$, and use Lemma 2.10, with $\xi = -(\varepsilon^{i_1}, \dots, \varepsilon^{i_k})^\top$, to calculate

$$\det(A(\varepsilon)_{\alpha\beta}) = \det \left(\mathbf{A} - \begin{pmatrix} \varepsilon^{i_1} & \dots & \varepsilon^{i_1} \\ & \dots & \\ \varepsilon^{i_k} & \dots & \varepsilon^{i_k} \end{pmatrix} \right) = \det(\mathbf{A}) + \sum_{l=1}^k \det([\mathbf{A} \mid_l \xi]) \quad (2.12)$$

Laplace determinant expansion along the l th column yields

$$\det([\mathbf{A} \mid_l \xi]) = - \sum_{h=1}^k (-1)^{l+h} \varepsilon^{i_h} \det(\mathbf{A}_{hl})$$

where \mathbf{A}_{hl} denotes the matrix obtained from \mathbf{A} by deleting the h th row and l th column. Hence we obtain that (2.12) equals

$$\det(\mathbf{A}) - \sum_{h=1}^k \varepsilon^{i_h} \sum_{l=1}^k (-1)^{l+h} \mathbf{A}_{hl} = \det(\mathbf{A}^\top) + \sum_{h=1}^k \varepsilon^{i_h} (-\det([\mathbf{A}^\top \mid_h \mathbf{1}]))$$

For small ε this expression (and with it the determinant in (2.12)) is positive (negative) if and only if the vector

$$[\det(\mathbf{A}^\top), -\det([\mathbf{A}^\top \mid_1 \mathbf{1}]), \dots, -\det([\mathbf{A}^\top \mid_k \mathbf{1}])]$$

is lexicopositive (lexiconegative). Hence we get

$$\text{sign}(\det(A(\varepsilon)_{\alpha\beta})) = \zeta(\mathbf{A}^\top)$$

which cannot be zero since the game is nondegenerate. The same calculation applied to $B(\varepsilon)$ yields the second half of the index calculation. \square

Proposition 2.9 makes it easy to calculate the index of any equilibrium component, and is the final ingredient for our definition of an index concept for extreme equilibria, the lex-index. According to Proposition 2.7, every extreme equilibrium contributes in a precise way to the equilibria of the perturbed game. This result, together with Proposition 2.9, allows for a suggestive definition of the lex-index of such an extreme equilibrium. More precisely, for an extreme equilibrium (x, y) of a game (A, B) , define $\mathcal{B}(x, y)$ to be the set of all lexicofeasible and complementary pairs of bases of (2.2), (2.3) that yield (x, y) as the corresponding basic solution. Now define the *lex-index* of (x, y) to be

$$\sum_{(\alpha, \beta) \in \mathcal{B}(x, y)} (-1)^{|\beta|+1} \zeta(A_{\alpha\beta}^\top) \zeta(B_{\alpha\beta}) \quad (2.13)$$

where $|\beta|$, $A_{\alpha\beta}$ and $B_{\alpha\beta}$ are defined as in Proposition 2.9. Essentially, the lex-index of an extreme equilibrium is the sum of the indices of all equilibria of the game $(A(\varepsilon), B(\varepsilon))$ that come from that extreme equilibrium. For an equilibrium in a nondegenerate game, the lex-index coincides with the usual definition of the index. We reformulate the results from Propositions 2.4, 2.7 and 2.9 in terms of the lex-index, and finally prove Theorem 2.1 from the introduction:

Corollary 2.11. *The index of an equilibrium component is the sum of the lex-indices over all extreme equilibria of this component.*

The concept of the lex-index for extreme equilibria can be nicely demonstrated using the following basic example, given by the matrices $A = \begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix}$, $B = \begin{pmatrix} 1 & 1 \\ 1 & 2 \end{pmatrix}$. This game has just one equilibrium component, whose index is $+1$. Its extreme equilibria are $(1, 0), (0, 1), (1, 0), (1, 0)$ and $(0, 1), (0, 1)$. The component consists of two maximal Nash subsets, i.e. maximal convex set of equilibria (Jansen, 1981). The two maximal Nash sets that form the equilibrium component in this game are given by the set $\{(1, 0), (p, 1 - p) \mid p \in [0, 1]\}$, the other by the set

$\{(q, 1 - q), (0, 1) \mid q \in [0, 1]\}$; they intersect in the extreme equilibrium $(1, 0), (0, 1)$. What are the indices of those extreme equilibria? It is quite obvious (by iteratively eliminating dominated strategies) that in the lexicographically perturbed game, all equilibria apart from $(0, 1), (0, 1)$ vanish. In terms of our algorithm, this corresponds to the observation that the game gives rise to just one lexicofeasible complementary pair of bases (in which the basic variables are u, y_2, r_1 and v, x_2, s_1). This means that the extreme equilibrium $(0, 1), (0, 1)$ has lex-index $+1$, while the other two extreme equilibria have lex-index zero.

Although in nondegenerate games the lex-index of an equilibrium coincides with its index, it still makes sense to distinguish between the two concepts for the following reason: The lex-index depends on the specific perturbation that we chose, which implies that, in general, it is not independent of the order of pure strategies in the chosen representation of the game. As an example, consider the degenerate 3×2 game

$$A = \begin{pmatrix} 6 & 0 \\ 5 & 2 \\ 3 & 3 \end{pmatrix}, B = \begin{pmatrix} 0 & 1 \\ 2 & 0 \\ 4 & 4 \end{pmatrix}$$

This game has one isolated equilibrium, $(2/3, 1/3, 0), (2/3, 1/3)$, and one equilibrium component whose extreme equilibria are $(0, 0, 1), (1/3, 2/3)$ and $(0, 0, 1), (0, 1)$. It is straightforward to see that the isolated equilibrium has index $+1$, which implies that the component has index 0 . In the perturbed game $(A(\varepsilon), B(\varepsilon))$, the equilibrium component decomposes into two equilibria, one of which is pure, hence has index $+1$. Since the indices of the two equilibria must add up to zero, the other equilibrium must have index -1 . We conclude that one of the extreme equilibria has lex-index $+1$, the other lex-index -1 .

Exchanging the two columns of the game essentially does not change the equilibrium structure; we get again an isolated equilibrium and an equilibrium component. This time, however, the two extreme equilibria of the component vanish if we perturb the game using our lexicographic perturbation. This implies that both of the extreme equilibria of the component have lex-index 0 .

To conclude this chapter, we summarize our results in the following algorithm for the computation of the index of the equilibrium component C of a degenerate bimatrix game (A, B) :

Algorithm 2.12. *Input:* An equilibrium component C of a bimatrix game (A, B) .

Output: The index of C .

Method: Enumerate all extreme equilibria of C , using, for example, one of the methods by Avis et al. (2010). For every extreme equilibrium, compute its lex-index as defined in (2.13), using the method $\text{lex-index}(x, y)$ in Figure 2.1 below. Take the sum of the lex-indices over all extreme equilibria of C ; this sum equals the index of C .

$\text{bases}(x, y)$:

$\mathcal{B} = \emptyset$;

$B(x) = \{\text{bases of (2.3) corresponding to } x\}$;

$B(y) = \{\text{bases of (2.2) corresponding to } y\}$;

for (α, β) in $B(x) \times B(y)$:

if (α, β) lexicofeasible and complementary, and $(\alpha, \beta) \notin \mathcal{B}$:

$\mathcal{B} = \mathcal{B} \cup \{(\alpha, \beta)\}$;

Output \mathcal{B} .

$\text{lex-index}(x, y)$:

$\mathcal{B}(x, y) = \text{bases}(x, y)$;

Compute

$$i(x, y) = \sum_{(\alpha, \beta) \in \mathcal{B}(x, y)} (-1)^{|\beta|+1} \zeta(A_{\alpha\beta}^\top) \zeta(B_{\alpha\beta})$$

where $A_{\alpha\beta}$, $B_{\alpha\beta}$ and $|\beta|$ are defined as in Proposition 2.9, and ζ is the sign function defined in Definition 2.8;

Output $i(x, y)$.

Figure 2.1: The lex-index method used in Algorithm 2.12, which in turn uses the bases method.

3

Equilibrium Tracing in Strategic-Form Games

3.1 Introduction

In this chapter we investigate several algorithms for the computation of Nash equilibria in strategic-form games. The algorithms by Lemke and Howson (1964) and van den Elzen and Talman (1991) for bimatrix games are complementary pivoting methods; both have been studied extensively. The difference between the two methods is that while the Lemke-Howson method only allows for a restricted (finite) set of paths, the van den Elzen-Talman algorithm can start at any mixed strategy pair, called prior, and hence allows to generate infinitely many paths. This implies that the van den Elzen-Talman algorithm is more flexible than the Lemke-Howson method. An even more versatile algorithm is the global Newton method by Govindan and Wilson (2003a), which works for finite strategic-form games. All three algorithms can be interpreted as homotopy methods, see Herings and Peeters (2010).

We investigate the relations between these three algorithms. We show that the Lemke-Howson and van den Elzen-Talman algorithms differ substantially: The van den Elzen-Talman algorithm, when started from a pure strategy and its best response as a prior, in general finds a different equilibrium than the corresponding Lemke-Howson method. This is not surprising since both algorithms can be un-

derstood as special cases of the global Newton method, but in very different ways. The Lemke-Howson algorithm has been shown to be a special case of the global Newton method in Govindan and Wilson (2003b); we prove the corresponding result for the van den Elzen-Talman algorithm. We generalize this observation to the statement that for N -player strategic-form games, the global Newton method implements the linear tracing procedure introduced by Harsanyi (1975).

As a special case of the global Newton method, the van den Elzen-Talman algorithm can generically find only equilibria of index $+1$. This leads us to the issue of traceability of equilibria. Following Hofbauer (2003), we call an equilibrium in a bimatrix game traceable if it is found by the van den Elzen-Talman algorithm from an open set of priors. As explained above, the van den Elzen-Talman algorithm allows for much greater flexibility than the Lemke-Howson method. Hence one might hope that, unlike the Lemke-Howson algorithm, it is powerful enough to find all equilibria of index $+1$. This raises the question if, generically, all equilibria of index $+1$ are traceable. We answer this question negatively by analyzing traceability in coordination games.

If a nondegenerate 3×3 coordination game has a completely mixed equilibrium, this equilibrium has index $+1$. In addition, the game has three pure strategy equilibria, also of index $+1$, and three equilibria of support size two, which have index -1 . Hofbauer (2003) noted that in a symmetric 3×3 coordination game, the completely mixed equilibrium (if it exists) is not traceable. We show that, in general, this is only correct as long as we restrict the starting points of the van den Elzen-Talman paths to symmetric strategy profiles. More precisely, we will see that the traceability of the completely mixed equilibrium in a 3×3 coordination game depends on the specific geometry of the best reply regions. We prove that for certain generic coordination games the completely mixed equilibrium is traceable. However, we also show that there is a generic set of coordination games whose completely mixed equilibrium is not traceable. Hence there is an open set in the space of 3×3 bimatrix games that all have an untraceable equilibrium of index $+1$. This implies that the flexibility of the van den Elzen-Talman algorithm does not ensure generic traceability of all equilibria of index $+1$.

This, in turn, has important consequences for the concept of sustainability. Myerson (1997) suggested to call an equilibrium sustainable if it can be reached by

Harsanyi's and Selten's tracing procedure from an open set of priors. Since the van den Elzen-Talman algorithm implements the linear tracing procedure, this notion of sustainability is for nondegenerate bimatrix games equivalent to the concept of traceability. Hence the results of this chapter imply that generically not all equilibria of index $+1$ will be sustainable.

The structure of this chapter is as follows: In Section 3.2 we give a short review of the van den Elzen-Talman method and analyze its relations to the Lemke-Howson algorithm. In Section 3.3, we give a brief exposition of the global Newton method, before showing that it encompasses the van den Elzen-Talman algorithm and, more generally, the linear tracing procedure, as a special case. Section 3.4 contains a discussion of traceability of equilibria.

A version of this chapter has been published in *Economic Theory* (Balthasar, 2010).

3.2 Van den Elzen-Talman versus Lemke-Howson

The *van den Elzen-Talman algorithm* was introduced by van den Elzen and Talman (1991). It is a homotopy method that finds an equilibrium by starting at an arbitrary prior and adjusting the players' replies.

Let (A, B) be an $m \times n$ bimatrix game. As usual, denote by Δ_m and Δ_n the strategy simplices of players one and two, respectively, and the strategy space by $\Delta := \Delta_m \times \Delta_n$. For a subset Z of a real vector space, and a real number α , denote by $\alpha \cdot Z$ the set $\{\alpha z \mid z \in Z\}$. Take an arbitrary *prior* or *starting point* $(\bar{x}, \bar{y}) \in \Delta$. For $t \in [0, 1]$, define a new game $(A, B)^t$, in which the players choose a strategy $x \in t \cdot \Delta_m$ and $y \in t \cdot \Delta_n$, respectively, and get the payoff given by the matrices A and B against the strategy profile

$$((1-t)\bar{x} + x, (1-t)\bar{y} + y)$$

The game $(A, B)^t$ thus is the game that we get from (A, B) when we restrict the strategy choices of the players to

$$\Delta^t := (1-t)(\bar{x}, \bar{y}) + t \cdot \Delta$$

The *van den Elzen-Talman algorithm* traces equilibria of the game $(A, B)^t$, starting at the prior (\bar{x}, \bar{y}) for $t = 0$, and reaching an equilibrium of (A, B) at $t = 1$. In

general, degeneracies can occur along the path even in nondegenerate games. A discussion on how to resolve these can be found in von Stengel et al. (2002).

The van den Elzen-Talman algorithm can also be described as a complementary pivoting procedure: A point $(x, y) \in t \cdot \Delta$ yields an equilibrium in the restricted strategy space Δ^t if and only if there are suitable vectors r, s and real numbers u, v such that the following equations and inequalities hold:

$$\begin{aligned}
 A \cdot ((1-t)\bar{y} + y) + r &= u\mathbf{1} \\
 B^\top \cdot ((1-t)\bar{x} + x) + s &= v\mathbf{1} \\
 x^\top \mathbf{1} = t, y^\top \mathbf{1} &= t \\
 x^\top r = 0, y^\top s &= 0 \\
 x, r, y, s &\geq 0
 \end{aligned} \tag{3.1}$$

The vectors x and y indicate how much weight is put on each strategy in addition to that given by $(1-t)\bar{x}$ and $(1-t)\bar{y}$. The slack variables r and s show how far from being optimal a strategy is against the other player's strategy. The real numbers u and v track the equilibrium payoff during the computation.

As a complementarity pivoting algorithm, the van den Elzen-Talman method can be understood as a special case of Lemke's algorithm (von Stengel et al., 2002). The latter is a method for solving linear complementarity problems, by augmenting the original problem by a new variable, whose coefficients are given by a so-called covering vector (Lemke, 1965). For the case of van den Elzen-Talman, the new variable is $1-t$. The covering vector is essentially given by the payoffs $A\bar{y}$ and $B^\top \bar{x}$ against the prior.

The van den Elzen-Talman algorithm can also be understood geometrically as a completely labelled path in the strategy space Δ . As usual, assume that the players' pure strategies are numbered $1, \dots, m$ for player one and $m+1, \dots, m+n$ for player two. Recall that the best reply region for a pure strategy j of player two is defined as

$$\Delta(j) = \{x \in \Delta_m \mid j \text{ is a best reply to } x\}$$

Now, for a point $p = (1-t) \cdot \bar{x} + t \cdot x \in \Delta_m$ define its *labels at time t* to be

$$\{j \mid p \in \Delta(j)\} \cup \{i \mid x_i = 0\}$$

and similarly for the other player. Then by (3.1), a point in the restricted strategy space Δ^t is an equilibrium of the game $(A, B)^t$ if and only if it is completely labelled at time t . The pivoting steps of the algorithm occur where one of the players picks up a new label, which then the other player can drop. An analogous description in terms of “picking up” and “dropping” labels can be used to describe the Lemke-Howson algorithm, see Section 1.3 or von Stengel (2002). For further details on the van den Elzen-Talman algorithm we refer the reader to Herings and Peeters (2010), von Stengel (2002), or the original papers by van den Elzen and Talman (1991, 1999).

For a nondegenerate bimatrix game, what happens in the van den Elzen-Talman algorithm if we take the prior \bar{x} to be any pure strategy vector and \bar{y} its unique best reply? This would correspond to a starting point of the Lemke-Howson algorithm, and one might expect the two algorithms to find the same equilibrium.

However, this is not true. An example where the van den Elzen-Talman and Lemke-Howson paths lead to different equilibria is given by the 3×3 bimatrix game

$$A = \begin{pmatrix} 4 & 4 & 4 \\ 0 & 0 & 6 \\ 5 & 0 & 0 \end{pmatrix}, B = \begin{pmatrix} 6 & 12 & 0 \\ 0 & 4 & 0 \\ 8 & 0 & 13 \end{pmatrix} \quad (3.2)$$

and starting point $\bar{x} = (0, 0, 1), \bar{y} = (0, 0, 1)$. We saw earlier that the Lemke-Howson algorithm from this starting point finds the equilibrium $(5/11, 0, 6/11), (4/5, 0, 1/5)$ (see example (1.1) and Figure 1.1). However, the van den Elzen-Talman algorithm starting at this prior finds the pure strategy equilibrium $(1, 0, 0), (0, 1, 0)$. This can be seen from a graphical description of the corresponding van den Elzen-Talman path, which we have given in Figure 3.1. Another interesting feature of this path is that the homotopy parameter shrinks at some point during the algorithm. A further discussion of the relationship between the two algorithms will be provided at the end of the next section.

3.3 Relationships to the global Newton method

The *global Newton method* for equilibrium computation was introduced by Govindan and Wilson (2003a); it is a homotopy method for the computation of Nash




Figure 3.1: The van den Elzen-Talman path from the prior $(0, 0, 1)$, $(0, 0, 1)$ for example (3.2). The left simplex is player one's, the right one player two's. Player one's strategies are labelled 1-3, player two's have labels 4-6. The labels in the simplex mark the players' best reply regions. The labels outside mark the edges of the simplex where the corresponding strategy is unplayed. The square dot is the equilibrium that is found by the Lemke-Howson algorithm, (see Figure 1.1 for a graphic description of the corresponding Lemke-Howson path). The black arrows give the path of the van den Elzen-Talman algorithm starting at $(0, 0, 1)$, $(0, 0, 1)$, and are numbered in the order in which they occur. The dashed lines trace the restricted strategy space Δ^t after step 5 (upper line) and step 7 (lower line).

equilibria in finite strategic-form games. For simplicity, we will first give a description of the algorithm for bimatrix games, and then explain how to generalize it to N -player games.

First we need to introduce a procedure of creating new games from old ones that goes back to the structure theorem by Kohlberg and Mertens (1986): Starting from an $m \times n$ bimatrix game (A, B) and directional (column) vectors $a \in \mathbb{R}^m$, $b \in \mathbb{R}^n$, define a new game $(A, B) \oplus (a, b)$ by adding the vector a to each column of A , and the vector b^\top to each row of B . Hence the game $(A, B) \oplus (a, b)$ is given by the matrices

$$A + \begin{pmatrix} a_1 & & a_1 \\ \vdots & \dots & \vdots \\ a_m & & a_m \end{pmatrix}, B + \begin{pmatrix} b_1 & \dots & b_n \\ & \vdots & \\ b_1 & \dots & b_n \end{pmatrix} \quad (3.3)$$

Note that in general this procedure changes the equilibria of the game.

The idea of the global Newton method is as follows: Assume we would like to calculate an equilibrium of a bimatrix game (A, B) . For any pair of directional vectors

(a, b) as above, consider the ray $\{(A, B) \oplus \lambda \cdot (a, b) \mid \lambda \geq 0\}$ in the space of games. Take the graph of the equilibrium correspondence (which is the correspondence that maps each game to the set of its equilibria) over that ray. The structure theorem by Kohlberg and Mertens (1986) implies that generically, i.e. for (a, b) outside a lower-dimensional set, this graph will be a semi-algebraic one-dimensional manifold with boundary, where boundary points are equilibria of the game (A, B) . If we can find an equilibrium for large λ and trace it along that manifold, we arrive at an equilibrium of the original game.

Although the idea is conceptually straightforward, its implementation is technically demanding. Govindan and Wilson take advantage of the differentiable structure which is implicit in the structure theorem. They convert the problem of tracing equilibria over a ray to one of calculating zeros of piecewise differentiable functions, and for this they use the “original” global Newton method due to Smale (1976) (hence the name of the method). For further details we refer the reader to the original paper by Govindan and Wilson (2003a).

For our purpose, all we need to know is that for a bimatrix game (A, B) and a pair of directional vectors (a, b) in a suitable Euclidean space, the global Newton method traces equilibria along the graph of the equilibrium correspondence over the ray

$$\{(A, B) \oplus \lambda \cdot (a, b) \mid \lambda \geq 0\}$$

In other words, for the graph \mathcal{E} of the equilibrium correspondence, the global Newton method traces equilibria along the set

$$\{((A, B) \oplus \lambda \cdot (a, b), (x, y)) \in \mathcal{E} \mid \lambda \geq 0\}$$

A crucial condition for the algorithm to work is that this set is nondegenerate, in the sense that it is a one-dimensional manifold (with boundary) without branch points (by this, we also mean that it may have no branch points “at infinity”). Generically, however, this nondegeneracy condition is satisfied.

The global Newton method can easily be extended to games with more than two players. Definition (3.3) means that for each player, a bonus is added to his payoff from each of his pure strategies, regardless of the other player’s strategy. This concept has an obvious extension to N -player strategic-form games Γ , where each player i has a “bonus vector” g_i . As in (3.3), we get a new game $\Gamma \oplus (g_1, \dots, g_N)$.

The global Newton method then traces equilibria over rays of the form

$$\{\Gamma \oplus \lambda \cdot (g_1, \dots, g_N) \mid \lambda \geq 0\}$$

We now prove that the global Newton method comprises the van den Elzen-Talman algorithm as a special case, and give a generalization of this result to the linear tracing procedure in the case of a finite strategic-form game. Let (A, B) be an $m \times n$ bimatrix game. Choose a starting point $(\bar{x}, \bar{y}) \in \Delta_m \times \Delta_n$. The van den Elzen-Talman algorithm traces the set

$$P_{ET}((A, B), (\bar{x}, \bar{y})) = \left\{ \begin{array}{l} (t, (x, y)) \in [0, 1] \times \Delta_m \times \Delta_n \mid (x, y) \in \Delta^t \text{ and } (x, y) \text{ is} \\ \text{an equilibrium of the game } (A, B)^t \end{array} \right\}$$

where Δ^t is the restricted strategy space and $(A, B)^t$ the corresponding game defined in Section 3.2.

For $\lambda \in \mathbb{R}$, define the game

$$(A, B)^{\bar{x}, \bar{y}}(\lambda) = (A, B) \oplus \lambda \cdot (A\bar{y}, B^\top \bar{x})$$

where \oplus is defined as in (3.3). Let \mathcal{E} be the graph of the equilibrium correspondence over the space of bimatrix games $\mathbb{R}^{m \times n} \times \mathbb{R}^{m \times n}$, and let

$$P_{GNM}((A, B), (\bar{x}, \bar{y})) = \{(A, B)^{\bar{x}, \bar{y}}(\lambda), (x, y) \in \mathcal{E} \mid \lambda \geq 0\}$$

be the set of equilibria over the ray of games $\{(A, B)^{\bar{x}, \bar{y}}(\lambda) \mid \lambda \geq 0\}$. This is the set traced by the global Newton method, when choosing as directional vector $(A\bar{y}, B^\top \bar{x})$. The following Proposition states that it is homeomorphic to the set $P_{ET}((A, B), (\bar{x}, \bar{y}))$, after removing the starting point $(0, (\bar{x}, \bar{y}))$ from the latter. This establishes the van den Elzen-Talman algorithm as a special case of the global Newton method.

Proposition 3.1. *Let (A, B) be an $m \times n$ bimatrix game. Choose a starting point $(\bar{x}, \bar{y}) \in \Delta_m \times \Delta_n$. Let $\lambda : (0, 1] \rightarrow \mathbb{R}^{\geq 0}$, $t \mapsto \frac{1}{t} - 1$. Then the map*

$$\begin{aligned} P_{ET}((A, B), (\bar{x}, \bar{y})) \setminus \{(0, (\bar{x}, \bar{y}))\} &\rightarrow P_{GNM}((A, B), (\bar{x}, \bar{y})) \\ (t, (1-t)\bar{x} + tx, (1-t)\bar{y} + ty) &\mapsto ((A, B)^{\bar{x}, \bar{y}}(\lambda(t)), (x, y)) \end{aligned}$$

is a homeomorphism.

Proof. In the game (A, B) , the payoff vector for player one against the strategy $(1-t) \cdot \bar{y} + t \cdot y$ for $y \in \Delta_n$ is

$$(1-t)A\bar{y} + tAy = ((1-t)(A\bar{y}, \dots, A\bar{y}) + tA)y$$

where we exploit the fact that $y^\top \mathbf{1} = 1$. Similarly the payoff vector for player two against a strategy $(1-t) \cdot \bar{x} + t \cdot x$ for $x \in \Delta_m$ is

$$\begin{aligned} (1-t)B^\top \bar{x} + tB^\top x &= ((1-t)(B^\top \bar{x}, \dots, B^\top \bar{x}) + tB^\top)x \\ &= ((1-t) \begin{pmatrix} \bar{x}^\top \cdot B \\ \vdots \\ \bar{x}^\top \cdot B \end{pmatrix} + tB)^\top x. \end{aligned}$$

Hence a strategy profile $((1-t) \cdot \bar{x} + t \cdot x, (1-t) \cdot \bar{y} + t \cdot y)$ in the restricted strategy space is an equilibrium of $(A, B)^t$ if and only if (x, y) is an equilibrium of the game $t \cdot (A, B) \oplus (1-t) \cdot (A\bar{y}, B^\top \bar{x})$.

Since the equilibria of a game remain unchanged by multiplication of the payoffs by a positive constant, we get that the set $P_{ET}((A, B), (\bar{x}, \bar{y}))$ is given by

$$\{(0, (\bar{x}, \bar{y}))\} \cup \left\{ (t, (1-t) \cdot \bar{x} + t \cdot x, (1-t) \cdot \bar{y} + t \cdot y) \mid t \in (0, 1], (x, y) \text{ is an equilibrium of the game } (A, B) \oplus \left(\frac{1}{t} - 1\right) \cdot (A\bar{y}, B^\top \bar{x}) \right\},$$

which ensures that our map maps indeed to $P_{GNM}((A, B), (\bar{x}, \bar{y}))$. Since it is obviously continuous, we just need to find a continuous inverse. This can be easily done by taking the inverse map to λ and taking the corresponding continuous map

$$P_{GNM}((A, B), (\bar{x}, \bar{y})) \rightarrow P_{ET}((A, B), (\bar{x}, \bar{y})) \setminus \{(0, (\bar{x}, \bar{y}))\}$$

□

The map from Proposition 3.1 can easily be extended to the point $(0, (\bar{x}, \bar{y}))$ by taking the one-point-compactification of $P_{GNM}((A, B), (\bar{x}, \bar{y}))$. As an immediate consequence we get that a van den Elzen-Talman path is a one-dimensional manifold (with boundary) without branch points if and only if the same holds for the corresponding path of the global Newton method (where a branch point at the starting point of the van den Elzen-Talman path would correspond to a branch point “at infinity” of the path of the global Newton method). If both paths satisfy this non-degeneracy condition, both algorithms will find the same equilibrium.

An alternative proof for Proposition 3.1 can be given using Lemke's algorithm: For bimatrix games, the global Newton method corresponds to Lemke's algorithm (with potentially non-positive covering vectors), see Govindan and Wilson (2003a). The bonus vector for the global Newton method essentially corresponds to the covering vector used in Lemke's algorithm. Using that the van den Elzen-Talman method is a special case of Lemke's algorithm, Proposition 3.1 follows.

Proposition 3.1 can be generalized to N -player games as follows: It has been proved in van den Elzen and Talman (1999) that the van den Elzen-Talman algorithm implements the linear tracing procedure, which was introduced by Harsanyi (1975). The linear tracing procedure is a method for selecting a Nash equilibrium in an N -player game; it plays a key role in the equilibrium selection theory developed by Harsanyi and Selten (1988). For any prior (i.e. mixed strategy profile), the linear tracing procedure traces equilibria over a set of games whose payoffs are given as a convex combination of the original payoffs, and payoffs against the prior. To make this more precise, choose an N -player strategic-form game Γ and a prior p , and denote by $\Gamma_i(\sigma)$ the payoff of player i against a mixed strategy combination σ . For $0 \leq t \leq 1$, define a game Γ^t , which has the same sets of players and strategies as Γ , but the payoff in Γ^t to player i from a strategy combination σ is defined as

$$\Gamma_i^t(\sigma) = t\Gamma_i(\sigma) + (1-t)\Gamma_i(\sigma_i, p_{-i}) \quad (3.4)$$

where (σ_i, p_{-i}) is the strategy combination that results from p by replacing player i 's strategy p_i by σ_i . The linear tracing procedure traces the graph of the equilibrium correspondence over the set of games $\{\Gamma^t \mid t \in [0, 1]\}$, which in almost all cases will be a one-dimensional manifold. For $t > 0$ we can divide the payoffs given in (3.4) by t without changing the equilibria of the game, and as in the proof above we can conclude the following generalization of Proposition 3.1, which is one of the central results of this chapter:

Theorem 3.2. *The global Newton method implements the linear tracing procedure.*

It has been proved in Govindan and Wilson (2003b) that the global Newton method also comprises the Lemke-Howson algorithm. Proposition 3.1 raises the question of how the latter algorithm as a special case of the global Newton method differs from the van den Elzen-Talman algorithm. If we take the i th unit vector

e_i for some pure strategy i of player one, the global Newton method for the ray $(A, B) \oplus \lambda \cdot (e_i, 0)$ corresponds to the Lemke-Howson algorithm with missing label i . An analogous statement holds for missing labels of player two; for further details see Govindan and Wilson (2003b). So the Lemke-Howson algorithm corresponds to taking unit vectors as directional vectors for the global Newton method, whereas the van den Elzen-Talman algorithm is based on directional vectors $(A\bar{y}, B^\top \bar{x})$. Further differences between the two algorithms will emerge in the analysis of coordination games in the next section: We will see that in this type of game, the Lemke-Howson algorithm only finds the pure strategy equilibria, whereas for certain coordination games, the van den Elzen-Talman method can also find the completely mixed equilibrium.

3.4 Traceability and the index of equilibria

In this section we discuss which equilibria can be traced by the van den Elzen-Talman algorithm. Of course every equilibrium can be found by taking it as starting point. However, we are only interested in those that are found generically. As suggested by Hofbauer (2003), we call an equilibrium of a nondegenerate bimatrix game *traceable* if it can be reached by the van den Elzen-Talman algorithm from an open set of priors. As explained in Section 3.1, traceability in this sense corresponds to a notion of sustainability suggested by Myerson (1997). Govindan and Wilson (2003a) have shown that, generically, every equilibrium found by the global Newton method has index $+1$. Proposition 3.1 then implies that only equilibria of index $+1$ are traceable.

The converse question is if, generically, every equilibrium of index $+1$ is traceable. This question has been discussed in Hofbauer (2003) in the context of sustainability. We answer it negatively by giving an analysis of *coordination games*. Following Hofbauer (2003), we define a coordination game to be a square bimatrix game (A, B) , where the matrices A and B have zeros on the diagonal and negative entries off the diagonal. We restrict our analysis to nondegenerate 3×3 coordination games. If such a game has a completely mixed equilibrium, this equilibrium has index $+1$, and will be our candidate for non-traceability. In addition, such a game has three pure strategy equilibria, also of index $+1$, and three equilibria of

support size two, which have index -1 .

It is straightforward that in such a game, the Lemke-Howson algorithm only finds the pure strategy equilibria. These equilibria are traced by the van den Elzen-Talman method as well, by starting from any prior nearby. However, compared to the Lemke-Howson method, the van den Elzen-Talman algorithm allows for a vast variety of starting points. The question is if this increased flexibility suffices to make the completely mixed equilibrium traceable as well. Hofbauer (2003) answered this question negatively for symmetric games, which in general is correct only as long as we restrict the van den Elzen-Talman algorithm to symmetric starting point. More precisely, in this section we show that the traceability of the completely mixed equilibrium depends on the type of coordination game at hand. On the one hand, we give a class of generic coordination games for which the completely mixed equilibrium is not traceable. This implies that the flexibility of the van den Elzen-Talman algorithm does not ensure generic traceability of all equilibria of positive index.

On the other hand, we prove that there are coordination games for which the completely mixed equilibrium can indeed be traced from an open set of starting points. Hence for this class of games, the van den Elzen-Talman algorithm is stronger than the Lemke-Howson method, in the sense that the equilibria found by the latter method are a proper subset of the traceable equilibria. This strengthens known observations that the van den Elzen-Talman algorithm in general finds more equilibria than the Lemke-Howson method. An example of an equilibrium found by the van den Elzen-Talman algorithm but not by the Lemke-Howson method has been given by van den Elzen and Talman (1991). However, the equilibrium considered in their example has negative index, hence is only found via non-generic van den Elzen-Talman paths and cannot be traceable.

We start our analysis of traceability in coordination games with the game given by

$$A = \begin{pmatrix} 0 & -1 & -1 \\ -1 & 0 & -1 \\ -1 & -1 & 0 \end{pmatrix} = B^\top \quad (3.5)$$

which corresponds to the “standard” coordination game usually given by the identity matrix. For this game, it is easy to see that the completely mixed equilibrium

is not traceable. The van den Elzen-Talman paths in this example are quite simple; as soon as a path arrives in the “same” best reply regions for both players, the corresponding pure strategy equilibrium is found, as depicted in Figure 3.2.




Figure 3.2: A van den Elzen-Talman path for example (3.5). The dots denote the prior, or starting point of the algorithm. The black arrows give the path of the algorithm; the dashed triangles trace the value of the restricted strategy space Δ^t after the first step of the algorithm.

Consider the following perturbation of the standard coordination game (3.5):

$$A = \begin{pmatrix} 0 & -1 & -c \\ -c & 0 & -1 \\ -1 & -c & 0 \end{pmatrix} = B^\top \quad (3.6)$$

where $c > 0$. We call such a game a c -coordination game. The edges between any two best reply regions for this game are given by the points $\frac{1}{1+c} \cdot (0, 1, c)$, $\frac{1}{1+c} \cdot (c, 0, 1)$ and $\frac{1}{1+c} \cdot (1, c, 0)$, each connected to $(1/3, 1/3, 1/3)$. We will prove that the slopes of those edges are crucial to whether the completely mixed equilibrium in such a game is traceable or not. For $1 < c < 2$, the edges are contained in the darkly shaded areas in Figure 3.3. This implies that the smaller angle between any of those edges and the boundary of the respective player’s strategy space is between 60° and 90° . This property is crucial in the proof of the following Theorem.

Theorem 3.3. *For any $1 < c < 2$, the completely mixed equilibrium in the corresponding 3×3 c -coordination game defined in (3.6) is not traceable. The same is true for any small (possibly non-symmetric) perturbation of such a game.*

We do not claim this result to be sharp. To the contrary, we would expect an




Figure 3.3: For a c -coordination game, the edges between two best reply regions are contained in the darkly shaded areas if $1 < c < 2$, and in the lightly shaded areas if $c > 2$.

analogous statement to be true for other suitable parameters c , for example $1/2 < c \leq 1$.

Proof. Fix a c -coordination game with $1 < c < 2$. We have to analyze the different paths that can be generated by the van den Elzen-Talman algorithm by choosing an arbitrary starting point (\bar{x}, \bar{y}) . First, observe that it suffices to only consider non-degenerate paths, i.e. paths without branch points or higher-dimensional degeneracies. This is essentially due to the fact that the van den Elzen-Talman algorithm can be understood as a special case of Lemke’s algorithm, as explained in Section 3.2. Hence van den Elzen-Talman paths can be seen as Lemke-paths, which for almost all covering vectors will be nondegenerate as long as the underlying linear complementarity problem is nondegenerate (Eaves, 1971). But since the coordination games at hand are nondegenerate, the corresponding linear complementarity problem is nondegenerate as well (von Stengel, 2002).

Our proof is best understood by geometrically following the paths generated from different starting points. For illustration we have done this for the last case in Figure 3.4. We will describe the pivoting steps of possible paths using the concept of “picking up” and “dropping” labels, as described in Section 3.2. Player one’s strategies are labelled 1-3, and player two’s have labels 4-6. Recall that a point $(1-t)\bar{x} + x$ of player one’s restricted strategy space has as labels firstly the best replies of player two against $(1-t)\bar{x} + x$, and secondly his own “unused” strategies, i.e. those i with $x_i = 0$. The labels for a point in the restricted strategy space for

player two are defined analogously. If we say that at some point a player picks up or drops a label, we always mean this to be a label in the restricted strategy space relevant at that point. Recall that $\Delta(i)$ denotes the i th best reply region. As explained above, our choice of $c \in (1, 2)$ means that the edges between any two best reply regions are in the darkly shaded areas depicted in Figure 3.3. This property, which we will refer to as Property (*) for the remainder of the proof, is essential for the geometric structure of the van den Elzen-Talman paths.

Now choose a starting point (\bar{x}, \bar{y}) which generates a nondegenerate van den Elzen-Talman path. We can assume without loss of generality that $\bar{x} \in \Delta(4)$. As always in a coordination game, if $\bar{y} \in \Delta(1)$, the equilibrium $(1, 0, 0), (1, 0, 0)$ is found straight away. Next, assume that $\bar{y} \in \Delta(2)$, and look at the different cases that may happen. The case of $\bar{y} \in \Delta(3)$ is symmetric (by rotating the strategy space and exchanging the two players), hence there is no need to discuss it.

In the first step of the algorithm, the homotopy parameter starts growing, while both players put weight on their respective best replies to the prior. Hence the first part of the van den Elzen-Talman path is given by $((1-t) \cdot \bar{x} + (0, t, 0), (1-t) \cdot \bar{y} + (t, 0, 0))$. This path is followed until a new label is picked up, i.e. until the path hits another best reply region.

- If it hits $\Delta(5)$ or $\Delta(1)$ first, the corresponding pure strategy equilibrium is found straight away, similarly as in Figure 3.2.
- If the path hits $\Delta(3)$ first, player two picks up label 3, which then player one can drop. This means that player one starts putting positive weight on his third strategy, while the homotopy parameter t needs to remain constant in order to enable the players to keep all necessary labels. Since by Property (*), the edge between $\Delta(4)$ and $\Delta(5)$ is steeper than the edge given by label 1 of player one's restricted strategy simplex, the path cannot hit $\Delta(5)$ during this process. If the path hits $\Delta(6)$, then the equilibrium $(0, 0, 1), (0, 0, 1)$ is found: Player two can drop label 6, which player one has just found, but needs to keep labels 2, 3 and 5. Since by Property (*), the edge given by label 5 of player two's restricted strategy simplex is contained in $\Delta(3)$, the only way this can happen is if the homotopy parameter t starts shrinking, until either of the players reaches the upper vertex of the restricted strategy sim-

plex Δ^t , where that player finds a label of his own. Assume this was player one, picking up label 2 (the case for player two works vice versa). Then player two can drop that label, i.e. leave the boundary between the best reply regions, and walk towards the upper vertex of his restricted strategy simplex (while the homotopy parameter stalls), until he picks up label 4. This in turn implies that now player one can also leave the boundary of the best reply regions. The homotopy parameter starts growing, while both players stay in the upper corner of their restricted strategy spaces, until the equilibrium $(0, 0, 1), (0, 0, 1)$ is found.

The only remaining case is for the first player's path to remain in $\Delta(4)$ until he reaches the upper vertex $(1-t) \cdot \bar{x} + (0, 0, t)$ of his restricted strategy simplex and picks up label 2. Then the homotopy parameter starts growing again until one of the following cases occur:

- (i) If the path hits $\Delta(1)$ first, then the equilibrium $(1, 0, 0), (1, 0, 0)$ is found, by an argument similar to Figure 3.2.
- (ii) If the path hits $\Delta(6)$ first, then the equilibrium $(0, 0, 1), (0, 0, 1)$ is found, again by a similar argument.
- (iii) If the path hits $\Delta(5)$ first, player one picks up label 5, which then player two can drop. Then the homotopy parameter stalls while player two puts more weight on his second strategy. At some point he arrives at $\Delta(2)$ again. Now player one can drop label 2, but needs to keep labels 1, 4 and 5. Since due to Property (*), the edge given by label 1 of player one's restricted strategy simplex is contained in $\Delta(5)$, the only way this can happen is if the homotopy parameter starts shrinking, while both players walk along the edge between the best reply regions they are on, away from the barycenter $(1/3, 1/3, 1/3)$. At some point, player one reaches the right vertex of his restricted strategy simplex, picking up label 3. This is bound to happen before player two reaches the left vertex of his restricted strategy simplex: Due to the history of the algorithm we can see that player two's relevant vertex is further away (in terms of the homotopy parameter) from the relevant boundary between best reply regions, than player one's. Player two can now drop label 3, i.e. leave $\Delta(3)$: While the homotopy parameter is stalling, he




Figure 3.4: A van den Elzen-Talman-path as in the proof of Theorem 3.3. The black line segments give the path of the algorithm, the dashed triangles trace the restricted strategy space Δ^t . The upper figure contains the first four steps of the algorithm, the lower one traces the whole path in greater detail.

walks towards the right vertex of his restricted strategy simplex, where he picks up label 4. Now player one can leave $\Delta(4)$, and the equilibrium $(0, 1, 0), (0, 1, 0)$ is found. For visualization, we have provided a graphic description of the last case in Figure 3.4.

So far we have proven that for any c -coordination game with $1 < c < 2$, the completely mixed equilibrium is not traceable. However, it is easy to verify that for any small (possibly non-symmetric) perturbation of such a c -coordination game, the arguments above are still valid. Hence we can extend our result to small perturbations of such c -coordination games. \square

As an immediate consequence, we get the central result of this section:

Corollary 3.4. *There is an open set in the space of 3×3 bimatrix games, such that every game in that set has an equilibrium of index +1 that is not traceable.*

We would like to conclude this section by proving that the mixed equilibrium of a c -coordination game as in (3.6) is traceable as soon as $c > 2$. This is essentially

due to the fact that for $c > 2$, the relevant angle of the edges between best reply regions and the boundary of the strategy space becomes smaller than 60° , i.e. those edges are contained in the lightly shaded areas of Figure 3.3, as opposed to the darkly shaded ones. This in turn implies that where the homotopy parameter t used to shrink before, now it starts growing, which enables us to find the completely mixed equilibrium.

Proposition 3.5. *For any $c > 2$, the completely mixed equilibrium in the corresponding 3×3 c -coordination game is traceable.*

Proof. Choose a prior $(\bar{x}, \bar{y}) \in \Delta(4) \times \Delta(2)$ such that \bar{x} is close to $\Delta(6)$ but far from $\Delta(5)$, while \bar{y} is closer to $\Delta(3)$ than \bar{x} is to $\Delta(5)$, but further away from $\Delta(3)$ than \bar{x} is from $\Delta(6)$. At the same time, let \bar{y} be close to $\Delta(1)$. “Close” and “far” are to be understood in terms of the homotopy parameter. It should become clear during the description of the envisaged path what precisely is needed.

Finally, let \bar{y} be such that the line from \bar{y} to $(1, 0, 0)$ intersects $\Delta(3)$. By these choices, we can generate the following van den Elzen-Talman path: The path hits $\Delta(3)$ first, where player two picks up label 3. This means that player one can put weight on his third strategy, while the homotopy parameter stalls, until his path hits $\Delta(6)$. Now, player two can drop label 6 but needs to keep labels 2, 3 and 5. Since $c > 2$, the edges between any two best reply regions have a different slope than for the case where $1 < c < 2$, as depicted in Figure 3.3. In particular, the edge between $\Delta(2)$ and $\Delta(3)$ is steeper than the edge given by label 5 of player two’s restricted strategy space. This implies that at this point of the algorithm, player two’s restricted strategy simplex is still contained in $\Delta(2)$. This means that unlike in the proof of Theorem 3.3, where in the analogous situation the homotopy parameter started shrinking, the homotopy parameter now needs to grow in order to enable player two to keep the necessary labels. Since player one needs to keep labels 1, 4 and 6, he moves towards the barycenter $(1/3, 1/3, 1/3)$ on the edge between $\Delta(4)$ and $\Delta(6)$, while player 2 moves away from $(1/3, 1/3, 1/3)$ on the edge between $\Delta(2)$ and $\Delta(3)$. At some point, player one arrives at $(1/3, 1/3, 1/3)$, where he reaches $\Delta(5)$ and picks up label 5. By our choice of \bar{y} “close” to $\Delta(1)$, we can assume that at this point, the barycenter $(1/3, 1/3, 1/3)$ is contained in player two’s restricted strategy space. Player two can now drop label 5, i.e. put positive weight on his second strategy, hence the homotopy parameter stalls and player two

moves along the edge between $\Delta(2)$ and $\Delta(3)$ until he, too, reaches $(1/3, 1/3, 1/3)$. Now in turn player one can put positive weight on his first strategy. This implies that the homotopy parameter can grow until it reaches 1, while both players remain at the completely mixed equilibrium. Quite obviously, this path can be generated from an open set of priors, hence the completely mixed equilibrium is traceable.

To conclude the proof, we give a numerical example of the path generated above: For $c = 3$, from the starting point given by $\bar{x} = (45/100, 35/100, 20/100)$, $\bar{y} = (15/100, 40/100, 45/100)$ the van den Elzen-Talman algorithm finds the completely mixed equilibrium of the corresponding c -coordination game, via a path as described above. \square

3.5 Open questions

The main open question that is raised by Theorem 3.3 and Corollary 3.4 is if similar results hold for the global Newton method: Is there an equilibrium of index +1 that is not found by the global Newton method, or more generally, is there an open set of games such that each of these games has an equilibrium of index +1 that is not traceable by the global Newton method? Due to Proposition 3.1, the latter result would imply our Theorem 3.3.

4

Index and Uniqueness of Symmetric Equilibria

4.1 Introduction

In this chapter we use constructions of polytopes to prove a theorem, conjectured by Hofbauer (2003), about symmetric Nash equilibria of symmetric two-player games. These games are important in evolutionary game theory, where a “mixed strategy” represents the frequencies of individual “pure strategies” that are played in a population.

In a strategic form game, a Nash equilibrium always exists but is not necessarily unique. An enormous literature in game theory (van Damme, 1987) considers concepts of equilibrium selection and refinement in order to suggest fewer, preferably unique, equilibria as “solutions” to a given game. Typically, equilibria are selected that are “stable” in some sense, for example under perturbations of the payoffs that define the game (see Kohlberg and Mertens, 1986, and the subsequent literature). Hofbauer (2003) discusses various desirable properties of “sustainable” equilibria, a concept suggested by Myerson (1997); often, these properties hold for the equilibria of index $+1$. For example, only equilibria of index $+1$ can be stable under some “Nash field”, that is, a vector field on the set of mixed strategy profiles whose rest points are the Nash equilibria.

Hofbauer (2003) conjectured that equilibria of index $+1$ are “potentially unique” in

the sense that there is an extended game, with additional strategies for the players, where the given equilibrium is unique. This has been proved for bimatrix games in von Schemde (2005), with a constructive geometric proof in von Schemde and von Stengel (2008). We prove the corresponding theorem for symmetric equilibria of symmetric two-player games (for definitions see Section 4.2):

Theorem 4.1. *For a nondegenerate symmetric $d \times d$ game (B^\top, B) , a symmetric equilibrium has symmetric index $+1$ if and only if it is the unique symmetric equilibrium in an extended symmetric game (G^\top, G) .*

As explained in Chapter 2, the index of an equilibrium is an involved topological notion. Theorem 4.1, however, characterizes the symmetric index in purely strategic terms, without resorting to any concepts from topology. For $d \times n$ bimatrix games (A, B) , the analogous statement to Theorem 4.1 holds with the word “symmetric” omitted. (In this chapter we use d rather than m for the number of rows of the game matrices because most of our geometric objects live in \mathbb{R}^d .) This result for bimatrix games was first proved in von Schemde (2005) using topological arguments. The statements for bimatrix and symmetric games are independent because for a symmetric game, the bimatrix game index may differ from the symmetric index (we give an example after Definition 4.3 below). Also, the symmetric game needs to be extended symmetrically, whereas in the bimatrix game setting, only strategies for one player are added.

We prove Theorem 4.1 using polytopes. The symmetric game given in Theorem 4.1 is used to define a simplicial d -polytope. The polytope is *labelled* in the sense that each vertex has a *label* in $\{1, \dots, d\}$. A facet is *completely labelled* if the set of labels of its d vertices is $\{1, \dots, d\}$. The completely labelled facets correspond to the symmetric equilibria of the game, and one “artificial equilibrium” associated with a special facet F_0 . The orientation of a completely labelled facet is equal to the index of the corresponding equilibrium, except for a change of sign in even dimension (see Lemma 4.7). According to a standard “parity argument” (Papadimitriou, 1994) known from the Lemke-Howson algorithm (Lemke and Howson, 1964; Shapley, 1974), completely labelled facets come in pairs of opposite orientation. We state this result in Proposition 4.9 below. Its proof uses a very intuitive geometric argument, which relies heavily on the simpliciality of the polytope.

The following is our central result in the polytope setting.

Theorem 4.2. *Let P^Δ be a labelled simplicial d -polytope with $\mathbf{0}$ in its interior, and let F_0 and F be two completely labelled facets of opposite orientation. Then there are labelled points such that the convex hull P_{ext}^Δ of these points and P^Δ has only F_0 and F as completely labelled facets.*

The added points in Theorem 4.2 will be used to define the added strategies in Theorem 4.1. In order to get from added points to added strategies, we introduce a special class of bimatrix games, which we call *unit vector games*. A unit vector game is a bimatrix game where the columns of the first player's payoff matrix are unit vectors. This concept generalizes the imitation games introduced by McLennan and Tourky (2007), where the first player's payoff matrix is the identity matrix. We show that each labelled simplicial polytope P^Δ corresponds to a unit vector game; the completely labelled facets of P^Δ correspond to the equilibria of this game (see Lemma 4.10). For the labelled polytope P^Δ defined from a symmetric game (B^\top, B) , the corresponding unit vector game is the imitation game (I, B) . Hence Lemma 4.10 generalizes the result from McLennan and Tourky (2007) that the symmetric equilibria of (B^\top, B) are in one-to-one correspondence to the equilibria of (I, B) .

Starting from a symmetric game (B^\top, B) , we use Theorem 4.2 to add points to the corresponding labelled polytope P^Δ . These added points can be used to extend the corresponding imitation game (I, B) to a unit vector game, by adding strategies for the column player. We then have to symmetrize this game, by adding suitable payoff rows. The added rows are essentially given by the payoffs of the first player in the extended game (see Lemma 4.11). This is the crucial step for deriving Theorem 4.1 from Theorem 4.2.

This Chapter is organized as follows: In Section 4.2 we give a short exposition of symmetric games and the symmetric index. Section 4.3 explains how symmetric games are linked to labelled polytopes. Symmetric equilibria of a symmetric game correspond to completely labelled facets of the underlying polytope. Section 4.4 introduces the natural concept of an orientation of such a completely labelled facet, which up to a dimension-dependent sign coincides with the symmetric index of the corresponding symmetric equilibrium. In Section 4.5 we introduce unit vector

games, and show how to use these to derive our main result, Theorem 4.1, from its geometric counterpart, Theorem 4.2.

In the subsequent sections, we give a constructive proof of Theorem 4.2, the main idea of which goes back to a proof sketch of a related result in the non-symmetric setting by von Schemde and von Stengel (2008). However, the authors omit crucial details, and their proof sketch relies on results about polytopes which we prove in this chapter. Section 4.6 describes a central concept for our construction, the “P-matrix prism”, which is a particular type of completely labelled polytope with only two completely labelled facets. A known result on P-matrices allows us to prove Theorem 4.2 using a “stack” of three P-matrix prisms. For this to work, we need to re-arrange the polytope P^Δ in a suitable way; this is done in Section 4.7. In Section 4.8 we use a stack of P-matrix prisms to prove Theorem 4.2 for the case that the two given completely labelled facets are disjoint. The non-disjoint case will be treated in Section 4.9. The final Section 4.10 mentions open problems.

This chapter is joint work with Bernhard von Stengel.

4.2 The symmetric index

A *symmetric game* is a bimatrix game (B^\top, B) for a square matrix B (denoting the payoffs to the column player), that is, the game remains unchanged if the players are exchanged. A *symmetric equilibrium* of a symmetric game is an equilibrium of the form (x, x) , where both players use the same mixed strategy. Any symmetric game has a symmetric equilibrium (Nash, 1951). A symmetric game may also have non-symmetric equilibria, but in certain situations - i.e. if the players have no way of determining which of the two possible player positions they are in - only the symmetric equilibria are considered. Symmetric games and equilibria have been studied in a variety of contexts, especially in evolutionary game theory (see, for example, Gale et al., 1950, or Hofbauer and Sigmund, 1998).

Savani and von Stengel (2006) introduced a symmetric version of the Lemke-Howson algorithm. Following Shapley (1974), this algorithm can be used to define a symmetric version of the index, as follows: Consider a nondegenerate symmetric game (B^\top, B) with $B > 0$, and a symmetric equilibrium (x, x) , and let B_{xx} be the

square matrix obtained from B by deleting all rows and columns not in the support of x . As in Section 1.3, nondegeneracy implies that this matrix B_{xx} has full rank. The symmetric index of a symmetric equilibrium can now be defined analogously to the “ordinary” index:

Definition 4.3. *Let (B^\top, B) be a nondegenerate symmetric game with $B > 0$ and let (x, x) be a symmetric equilibrium with $k = |\text{supp}(x)|$. The symmetric index of (x, x) is defined as*

$$(-1)^{k+1} \text{sign det}(B_{xx}) \tag{4.1}$$

In analogy to the non-symmetric case, the symmetric version of the index has a straightforward interpretation in terms of the symmetric Lemke-Howson algorithm: Symmetric equilibria at opposite ends of Lemke-Howson paths have opposite symmetric index (Garcia and Zangwill, 1981). Like for the “ordinary” index, there are multiple ways of defining the symmetric index; for example, a version of the symmetric index based on the Poincaré index of the replicator dynamics has been suggested in Hofbauer and Sigmund (1998).

In a symmetric bimatrix game, the “ordinary” index (as in Definition 4.3) is in general different from the symmetric index. For example, the symmetric 2×2 game (A, B) of “chicken”, where $A = B = [e_2 \ e_1]$, has two non-symmetric pure equilibria and a mixed equilibrium which is the only symmetric equilibrium. That mixed equilibrium has index -1 for the bimatrix game, but symmetric index $+1$.

The symmetric index has the following properties, which require that its sign alternates with the parity of the support size as in (4.1) (compare Proposition 2.3 or (von Schemde and von Stengel, 2008, Proposition 2) for the corresponding statement for the “ordinary” index).

Proposition 4.4. *In a nondegenerate symmetric game, the symmetric index of a symmetric equilibrium*

- (a) *is $+1$ or -1 ;*
- (b) *does not change when adding a positive constant to all payoffs;*
- (c) *only depends on the payoffs in the support of the symmetric equilibrium;*
- (d) *does not depend on the order of the players’ pure strategies;*
- (e) *is $+1$ for any pure-strategy symmetric equilibrium; and*

(f) *the sum of the symmetric indices over all symmetric equilibria is +1.*

The proof of this proposition works analogously to the proof of the corresponding properties of the “ordinary index” (see, for example, von Schemde and von Stengel, 2008). As mentioned, (a) holds because the game is nondegenerate and B_{xx} has therefore full rank. It is easy to see that for a square matrix C and scalar s , the determinant $\det(C + sE)$ is a linear function of s , which is not constant if C is nonsingular, and does not change sign if $s \geq 0$ and $C > 0$. This implies (b), and it is the reason why we require that $B > 0$ in Definition 4.3.

In Proposition 4.4, claim (c) holds by definition, and shows that the index does not change when considering the equilibrium in an extended game with additional strategies (which are not played in the equilibrium). Condition (d) holds because rows and columns are exchanged equally to maintain the symmetry of the game. Property (e) is desirable, as discussed in Myerson (1997) and Hofbauer (2003), because pure-strategy equilibria are particularly convincing solutions to a game. Property (f) follows from a “parity argument” that we will prove in Proposition 4.9. It implies that a unique symmetric equilibrium must have index +1.

An example is the “coordination game” with the $d \times d$ identity matrix I as payoff matrix (its payoffs are only nonnegative, and not all positive, but Definition 4.3 still applies). In this game, any nonempty set $S \subseteq \{1, \dots, d\}$ of pure strategies defines a symmetric equilibrium (x, x) with $\text{supp}(x) = S$ and x as the uniform distribution on S . Its symmetric index is +1 if $|S|$ is odd, otherwise -1 .

4.3 Polytopes and symmetric equilibria

Polyhedra have been used since Vorob’ev (1958) to represent equilibria of bimatrix games. Consider a nondegenerate $d \times n$ bimatrix game with payoff matrix $B = [b_1 \cdots b_n]$ for player 2. (A special case is a symmetric $d \times d$ game (B^\top, B) with $n = d$.) Assume that the polyhedron

$$P = \{x \in \mathbb{R}^d \mid x \geq \mathbf{0}, x^\top B \leq \mathbf{1}^\top\} \quad (4.2)$$

is bounded. Recall that P is called a *best reply* polytope. Any x in $P - \{\mathbf{0}\}$ is interpreted as a mixed strategy $x/\mathbf{1}^\top x$ of player 1. Any tight inequality $x^\top b_j \leq 1$

describes a pure best reply j against x with payoff $1/\mathbf{1}^\top x$ (because any column other than j gives at most that payoff).

We require that P is bounded for the following reason (see also von Stengel (2002, Fig. 2.5) for a more detailed geometric interpretation): P is bounded if and only if the function $x \mapsto \mathbf{1}^\top x$ is bounded on P . Equivalently, player 2's best-reply payoff to any mixed strategy x is always positive, because $x^\top b_j \leq 0$ for all columns b_j of B would imply that $x \cdot t \in P$ for arbitrarily large scalars $t > 0$. Clearly, P is bounded if $B > 0$, but possibly also when some entries of B are negative, as in the following example:

$$B = \begin{pmatrix} 1/2 & -1/3 \\ -1/4 & 1 \end{pmatrix}, \quad P = \text{conv} \left\{ \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 2 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \begin{pmatrix} 3 \\ 2 \end{pmatrix} \right\} \quad (4.3)$$

Here, the three vertices of P other than $\mathbf{0}$ represent the pure strategies of player 1 and the mixed strategy $(3/5, 2/5)^\top$. For the symmetric game (B^\top, B) , these three vertices define symmetric equilibria, the pure ones with payoffs $1/2$ and 1 , respectively, and the mixed one with payoff $1/5$.

We will always start with a matrix $B > 0$, but will later add inequalities $x^\top b \leq 1$ to P with vectors b that may have negative entries. This is allowed because then P stays bounded. The columns b , and suitable rows, will be added to the payoff matrix to obtain an extended game.

Because the given $d \times n$ game is nondegenerate, no mixed strategy x of player 1 has more than $|\text{supp}(x)|$ pure best replies. Equivalently, no more than d inequalities in (4.2) are tight for any x in P . This means that P is a simple polytope, and that none of the inequalities $x^\top b_j$ is redundant in the sense that it can be omitted without changing the polytope (except when it defines the empty face, which means that the respective pure strategy j is never a best reply; then the j th column can be omitted from the game). We assume that the game is nondegenerate, so P is simple, and each binding inequality of P in (4.2) defines a facet.

Recall that each facet of P gets a *label* as follows. For $1 \leq i \leq d$, the facet $\{x \in P \mid x_i = \mathbf{0}\}$ gets label i . For $1 \leq j \leq n$, the facet $\{x \in P \mid x^\top b_j = 1\}$ gets label j . Any point x in P has the labels of the facets it lies on.

Consider a symmetric game (B^\top, B) , where $n = d$, and a point x in $P - \{\mathbf{0}\}$, which corresponds to a mixed strategy $x' = x/\mathbf{1}^\top x$. Then a label i of x represents either

an unplayed strategy of player 1 (when $x_i = 0$), or a best reply of player 2 ($x^\top b_i = 1$). Hence, by the best reply condition, the mixed strategy pair (x', x') defines a symmetric equilibrium if and only if x is completely labelled.

The vertex $\mathbf{0}$ of P is completely labelled, but it does not define an equilibrium. However, $\mathbf{0}$ serves as a starting point for a symmetric version of the Lemke-Howson algorithm (Savani and von Stengel, 2006, p. 402). The symmetric Lemke-Howson algorithm computes a path of edges of P which starts at a completely labelled vertex of P , for example $\mathbf{0}$, and ends at a different completely labelled vertex. The endpoints of any Lemke-Howson path have opposite symmetric index, where $\mathbf{0}$ has index -1 in agreement with (4.1) when $k = 0$. This implies Proposition 4.4(f). We prove a dual version of this observation in Proposition 4.9.

We will use the *polar* (or dual) polytope P^Δ instead of P . Suppose Q is a polytope,

$$Q = \{x \in \mathbb{R}^d \mid x^\top c_i \leq 1, 1 \leq i \leq k\} \quad (4.4)$$

with vectors c_1, \dots, c_k in \mathbb{R}^d . Then the *polar* (Ziegler, 1995) of Q is given by

$$Q^\Delta = \text{conv}\{c_1, \dots, c_k\} \quad (4.5)$$

The polytope P in (4.2) has to be translated in order to have $\mathbf{0}$ in its interior so that it can be written in the form (4.4). Moreover, it is convenient to have the negative unit vectors $-e_i$ as vertices of P^Δ , by translating P to the polytope $P' = \{x - \mathbf{1} \mid x \in P\}$. Then $\mathbf{0}$ is in the interior of P' if $\mathbf{1}$ is in the interior of P (like in the example (4.3)), that is, if

$$\mathbf{1}^\top b_j < 1 \quad \text{for } 1 \leq j \leq n \quad (4.6)$$

This can be assumed without loss of generality by multiplying all payoffs in B with a suitably small positive constant, which does not change the game.

Then $x' \in P' = \{x - \mathbf{1} \mid x \in P\}$ if and only if $x' + \mathbf{1} \geq \mathbf{0}$ and $(x' + \mathbf{1})^\top B \leq \mathbf{1}$, that is, $-x'_i \leq 1$ for $1 \leq i \leq d$ and $x'^\top b_j / (1 - \mathbf{1}^\top b_j) \leq 1$ for $1 \leq j \leq n$. Writing P^Δ instead of P'^Δ , we therefore obtain

$$P^\Delta = \text{conv}(\{-e_1, \dots, -e_d\} \cup \{b_j / (1 - \mathbf{1}^\top b_j) \mid 1 \leq j \leq n\}) \quad (4.7)$$

The facets of P^Δ correspond to the vertices of P and vice versa (Ziegler, 1995). The polytope P^Δ is simplicial (i.e. every facet has exactly d vertices) because P

is simple. The labels of the facets of P become labels of the vertices of P^Δ . By construction, these vertices are $-e_i$ with label i for $1 \leq i \leq d$, and $b_j/(1 - \mathbf{1}^\top b_j)$ with label j for $1 \leq j \leq n$.

The facet corresponding to the vertex $\mathbf{0}$ of P is given by

$$F_0 = \text{conv}\{-e_1, \dots, -e_d\} \quad (4.8)$$

Because $F_0 = \{x \in P^\Delta \mid -\mathbf{1}^\top x = 1\}$, the normal vector of F_0 is $-\mathbf{1}$, which is the vertex of P' that is the translated vertex $\mathbf{0}$ of P .

In general, a facet F of P^Δ has normal vector v if $F = \{x \in P^\Delta \mid v^\top x = 1\}$ and $v^\top x \leq 1$ is valid for all x in P^Δ . The normal vectors of facets F other than F_0 represent mixed strategies, as follows.

Lemma 4.5. *Let $F \neq F_0$ be a facet of P^Δ in (4.7) with normal vector v . Then v represents the mixed strategy $x = (v + \mathbf{1})/\mathbf{1}^\top(v + \mathbf{1})$, and $x_i = 0$ if and only if $-e_i \in F$ for $1 \leq i \leq d$. Any other label j of F , so that $b_j/(1 - \mathbf{1}^\top b_j)$ is a vertex of F , represents a pure best reply to x .*

Proof. This holds because the polar of the polar is the original polytope (Ziegler, 1995). More precisely, $P^{\Delta\Delta}$ is P' above, so the normal vector $v = (v_1, \dots, v_d)^\top$ is a vertex of P' and thus $v + \mathbf{1}$ is a vertex of P in (4.2). If $-e_i \in F$, then $v_i = -1$ and therefore $x_i = 0$, and vice versa. \square

Lemma 4.5 means that the labels of a facet F of P^Δ , whose normal vector represents a mixed strategy x , are the unplayed pure strategies in x or the pure best replies to x . Observe that by nondegeneracy, every symmetric equilibrium (x, x) of (B^\top, B) gives rise to a facet of P^Δ whose normal is $x - \mathbf{1}$, suitably scaled. Together with the following result, this implies that the symmetric equilibria of (B^\top, B) are in one-to-one correspondence with the completely labelled facets of P^Δ :

Corollary 4.6. *Let $F \neq F_0$ be a facet of P^Δ in (4.7) with normal vector v , and let $x = (v + \mathbf{1})/\mathbf{1}^\top(v + \mathbf{1})$ as in Lemma 4.5. Then (x, x) is a symmetric equilibrium of (B^\top, B) if and only if F has all labels $1, \dots, d$.*

Proof. An equilibrium (x, x) is given by those mixed strategies x so that for all $i = 1, \dots, d$, either $x_i = 0$ (that is, $-e_i$ with label i is a vertex of F) or i is a best reply to x . By Lemma 4.5, this means F has all labels. \square

A more direct representation of mixed strategies as normal vectors of facets has been considered in Bárány et al. (2005). The representation uses an unbounded polyhedron rather than a polytope, as follows. Consider the columns b_j of B as points in \mathbb{R}^d and let the polyhedron \hat{P} be the *nonnegative convex hull* of these points; that is, \hat{P} is the intersection of all halfspaces with nonnegative normal vectors that contain all points b_1, \dots, b_n . Then a normal vector x of any facet of \hat{P} , is, suitably scaled, a mixed strategy where the points b_j that lie on the facet are the pure best replies j to x . By construction, the polytope P^Δ is, with the exception of the additional facet F_0 , combinatorially equivalent to the polyhedron \hat{P} . The vertices b_j of \hat{P} are scaled to become vertices of P^Δ in (4.7).

We will later enlarge P^Δ by adding points c so that (among other things) F_0 remains a facet of $\text{conv}(P^\Delta \cup \{c\})$, that is, $-\mathbf{1}^\top c < 1$. These points correspond to additional columns b of the game matrix B given by

$$b = c / (1 + \mathbf{1}^\top c) \tag{4.9}$$

because then $b / (1 - \mathbf{1}^\top b) = c$ in agreement with (4.7).

4.4 Oriented facets

Considering the simplicial polytope P^Δ in (4.7) rather than the simple polytope P in (4.2) has the advantage that orientations of facets are easily defined and visualized. If F is a completely labelled facet of P^Δ , we assume that its vertices a_1, \dots, a_d are given in the order of their labels, that is, a_i has label i , for $1 \leq i \leq d$. Then the *orientation* of F is

$$\text{sign det}[a_1 \cdots a_d] \tag{4.10}$$

The orientation of F coincides with the symmetric index of the corresponding symmetric equilibrium, except for a change of sign in even dimension:

Lemma 4.7. *Let (B^\top, B) be a nondegenerate symmetric $d \times d$ game, and let P^Δ be the polytope in (4.7) with $n = d$. Then the orientation of a completely labelled facet F of P^Δ , multiplied by $(-1)^{d+1}$, is the symmetric index of the corresponding symmetric equilibrium, where F_0 corresponds to the artificial equilibrium with index -1 .*

Chapter 4. Index and Uniqueness of Symmetric Equilibria

Proof. When $F = F_0$, then its orientation is $\text{sign det}(-I)$ with the negative identity matrix $-I$, which is $+1$ when d is even and -1 when d is odd.

Let (x, x) be a symmetric equilibrium of (B^\top, B) , and let $F = \text{conv}\{a_1, \dots, a_d\}$ be the corresponding completely labelled facet of P^Δ . Here, a_i has label i for $1 \leq i \leq d$, where $a_i = -e_i$ if $i \notin \text{supp}(x)$ by Lemma 4.5, and $a_j = b_j / (1 - \mathbf{1}^\top b_j)$ for $j \in \text{supp}(x)$. For the sign of the determinant, we can ignore the scalar $1 / (1 - \mathbf{1}^\top b_j)$, which is positive by (4.6). Hence, with $k = |\text{supp}(x)|$,

$$\text{sign det}[a_1, \dots, a_d] = \text{sign}(-1)^{d-k} \det(B_{xx}) = (-1)^{d+1} (-1)^{k+1} \text{sign det}(B_{xx})$$

which proves the claim. \square

In general, the orientation of a nonsingular matrix is the sign of its determinant. When a facet is not completely labelled, there is no natural order of writing down its vertices as columns of a matrix. However, two *adjacent* facets share $d - 1$ vertices, so by keeping these $d - 1$ columns fixed, the respective matrices differ in only one column. The following lemma states that these matrices have opposite orientation. It is very intuitive in low dimension, which suggests its straightforward proof.

Lemma 4.8. *Consider a simplicial d -polytope with $\mathbf{0}$ in its interior, and two adjacent facets with vertices b, a_2, \dots, a_d and c, a_2, \dots, a_d , respectively. Then $[b \ a_2 \ \dots \ a_d]$ and $[c \ a_2 \ \dots \ a_d]$ have opposite orientation.*

Proof. We show that there are positive reals s and t so that $bs + ct$ is in the linear span of a_2, \dots, a_d . Then

$$\mathbf{0} = \det[bs + ct \ a_2 \ \dots \ a_d] = s \cdot \det[b \ a_2 \ \dots \ a_d] + t \cdot \det[c \ a_2 \ \dots \ a_d]$$

which implies that $\det[b \ a_2 \ \dots \ a_d]$ and $\det[c \ a_2 \ \dots \ a_d]$ have opposite sign as claimed; the determinants are nonzero because the hyperplanes through the two facets do not contain $\mathbf{0}$.

Let s, t , and r_2, \dots, r_d be reals, not all zero, so that

$$bs + ct + \sum_{i=2}^d a_i r_i = \mathbf{0} \tag{4.11}$$

where clearly $s \neq 0, t \neq 0$, and w.l.o.g. $s > 0$. Let v and w be the normal vectors to the two facets, so that $v^\top a_i = 1, v^\top b = 1, v^\top c < 1$ and $w^\top a_i = 1, w^\top c = 1, w^\top b < 1$,

for $2 \leq i \leq d$. Then multiplying (4.11) with both v^\top and w^\top gives $v^\top(bs + ct) = w^\top(bs + ct)$, that is, $(1 - w^\top b)s = (1 - v^\top c)t$ and thus $t = (1 - w^\top b)s / (1 - v^\top c) > 0$ where by (4.11) $bs + ct$ is in the linear span of a_2, \dots, a_d . \square

The following observation is proved with the path-following algorithm of Lemke and Howson (1964), in its symmetric form (Savani and von Stengel, 2006, p. 402), applied to the polar polytope of P in (4.2). It is the classic form of a ‘‘polynomial parity argument with direction’’ that defines the computational class PPAD (Papadimitriou, 1994). It is similar to the well-known proof by Cohen (1967) of Sperner’s Lemma (Sperner, 1928). We give a simplified version of the proof by Shapley (1974) that is based on exchanging columns of determinants. Lemke and Grotzinger (1976) give a similar proof based on abstract orientations of simplices in an oriented pseudo-manifold (see also Eaves and Scarf (1976), and Todd (1976), for orientation and index methods in the context of simplicial concepts).

Proposition 4.9. *In a labelled simplicial d -polytope with $\mathbf{0}$ in its interior, the completely labelled facets come in pairs of opposite orientation.*

Proof. Consider all facets that have all labels except possibly label 1. This includes any completely labelled facet, which we write as a matrix $[a_1 \cdots a_d]$ where vertex a_i of the facet has label i , for $1 \leq i \leq d$. The other facets have vertices $b_i, a_2, \dots, a_i, \dots, a_d$, where a_j has label j for $2 \leq j \leq n$ and b_i has the duplicate label $i \in \{2, \dots, n\}$. For these facets, we consider the two matrices

$$[b_i \ a_2 \ \cdots \ a_i \ \cdots \ a_d] \quad \text{and} \quad [a_i \ a_2 \ \cdots \ a_{i-1} \ b_i \ a_{i+1} \ \cdots \ a_d] \quad (4.12)$$

which have determinants of opposite sign. Consider all these matrices as nodes of a bipartite graph, with the matrices of negative determinant in one partition class and those of positive determinant in the other. Connect the two matrices in (4.12) by a ‘‘blue’’ edge. Secondly, join any other two matrices by a ‘‘red’’ edge if they have the same last $d - 1$ columns. This defines two adjacent facets. Their common $d - 1$ columns are not contained in any other matrix. By Lemma 4.8, the two matrices have opposite orientation, so the graph is indeed bipartite.

Every node in that graph has degree one or two. Any such graph is a collection of paths and cycles. The nodes of degree one, which are the endpoints of the paths, correspond to completely labelled facets and are only incident to a red edge. The

other nodes are also incident to a blue edge. Any path starts and ends with a red edge and is of odd length because the colors of the edges on the path alternate. Hence, the endpoints of any path have opposite orientation, as claimed. \square




Figure 4.1: Example of the path-following argument used to show Proposition 4.9. A vertex a_i or b_i has label i .

Figure 4.1 illustrates the preceding proof for $d = 3$. Each vertex a_i or b_i has label i in $\{1, 2, 3\}$. There is only a single path, which we start at $[a_1 \ a_2 \ a_3]$ which corresponds to a completely labelled facet with negative orientation (clockwise order of a_1, a_2, a_3). That path ends at $[b_1 \ a_2 \ b_3]$, oriented positively (anticlockwise). The path corresponds to the following sequence of matrices, with alternating orientation, and red edges shown with “ \rightarrow ”:
 $[a_1 \ a_2 \ a_3] \rightarrow [b_2 \ a_2 \ a_3], [a_2 \ b_2 \ a_3] \rightarrow [b_3 \ b_2 \ a_3], [a_3 \ b_2 \ b_3] \rightarrow [a_2 \ b_2 \ b_3], [b_2 \ a_2 \ b_3] \rightarrow [b_1 \ a_2 \ b_3]$.
 Only “red” edges are shown in Figure 4.1, as arrows from a matrix with negative to one with positive orientation, so a facet as in (4.12) where label 1 is missing is both endpoint and starting point of a red arrow; the “blue” edges just refer to an exchange of matrix columns.

The change from one simplicial facet to another is equivalent to a pivoting step as used in the simplex algorithm (Dantzig, 1963). The described paths may be exponentially long (Morris, 1994).

A dual best reply polytope P^Δ has the completely labelled facet F_0 , and therefore by Proposition 4.9 at least one other completely labelled facet F . This implies that every nondegenerate symmetric game has at least one symmetric equilibrium. (This is also true for degenerate games, with additional considerations.)

4.5 From unit vector to symmetric games

In this section, we show how Theorem 4.2 implies Theorem 4.1. Given a non-degenerate $d \times d$ symmetric game (B^\top, B) , we consider the labelled polytope P^Δ in (4.7). Then the symmetric equilibrium of index +1 considered in Theorem 4.1 corresponds to a completely labelled facet F which has opposite orientation to F_0 . According to Theorem 4.2, there are additional points c in \mathbb{R}^d so that the convex hull P_{ext}^Δ of these points and of P^Δ has only F and F_0 as completely labelled facets. Using (4.9), the added points correspond to added columns of the matrix B which then has n columns for some $n \geq d$. Furthermore, these points have labels in $\{1, \dots, d\}$. The following lemma shows how these labels can be used to define a bimatrix game whose equilibria correspond to the completely labelled facets of the extended polytope P_{ext}^Δ .

Lemma 4.10. *Consider a labelled simplicial d -polytope Q with $\mathbf{0}$ in its interior, spanned by a set of vertices*

$$\{-e_1, \dots, -e_d, c_1, \dots, c_n\}$$

so that F_0 in (4.8) is a facet of Q . Let $-e_i$ have label i for $1 \leq i \leq d$, and let c_j have label $l(j) \in \{1, \dots, d\}$ for $1 \leq j \leq n$. Let (U, B) be the $d \times n$ bimatrix game with $U = [e_{l(1)} \cdots e_{l(n)}]$ and $B = [b_1 \cdots b_n]$, where $b_j = c_j / (\mathbf{1}^\top c_j)$ for $1 \leq j \leq n$. Then the completely labelled facets F of Q , with the exception of F_0 , are in one-to-one correspondence to the Nash equilibria (x, y) of the game (U, B) , where x corresponds to F as in Lemma 4.5, and y is a suitable unit distribution.

Proof. Consider a facet F of Q so that $F = \text{conv}(\{-e_i \mid i \in K\} \cup \{c_j \mid j \in J\}) \neq F_0$. Let v be the normal vector to F , and let $x = (v + \mathbf{1}) / \mathbf{1}^\top (v + \mathbf{1})$ as in Lemma 4.5. Then x is a mixed strategy of player 1, which has support $\{1, \dots, d\} - K$, a set of size $|J|$. Furthermore, J is the set of pure best replies to x by player 2, who has payoff matrix B .

In order to obtain an equilibrium (x, y) of (U, B) for some mixed strategy $y \in \mathbb{R}^n$ of player 2, only best replies may be played with positive probability x_i or y_j . For player 2, this means that $y_j > 0$ only if $j \in J$. For player 1, we need $|J|$ pure best replies. Because the columns of player 1's payoff matrix U are unit vectors, this works only if $y_j = 1/|J|$ for $j \in J$ and $y_j = 0$ otherwise, and if for every $i \in$

$\text{supp}(x) = \{1, \dots, d\} - K$, there is some $j \in J$ so that $i = l(j)$ because then column j of U is the unit vector $e_{l(j)}$. This is exactly the condition that the set of labels of F , namely $K \cup \{l(j) \mid j \in J\}$, is $\{1, \dots, d\}$, that is, F is completely labelled, as claimed. \square

The bimatrix game (U, B) considered in the previous lemma may be called a *unit vector game*, that is, all columns of player 1's payoff matrix U are suitable unit vectors. A special case is an *imitation game* (I, B) where B is a square matrix and I is the identity matrix. Imitation games have been introduced by McLennan and Tourky (2007), who showed that the symmetric equilibria (x, x) of a symmetric game (B^\top, B) are in one-to-one correspondence with the equilibria (x, y) of the imitation game (I, B) , where y is the uniform distribution on $\text{supp}(x)$. This observation can be used to apply computational hardness results about symmetric games to bimatrix games. Furthermore, the symmetric index of (x, x) is equal to the index of (x, y) in the bimatrix game (I, B) .

The following lemma provides the main step for deriving Theorem 4.1 from Theorem 4.2. It explains how to get from a particular extension of an imitation game (I, B) to a symmetric extension of the corresponding symmetric game (B^\top, B) .

Lemma 4.11. *Consider $d \times d$ matrices I and B , where I is the identity matrix, and $d \times k$ matrices U and B' , where all columns of U are unit vectors, and let*

$$G = \begin{pmatrix} B & B' \\ U^\top & 0 \end{pmatrix} \quad (4.13)$$

Then any symmetric equilibrium (z, z) of (G^\top, G) gives rise to a Nash equilibrium (x, y) of the unit vector game $([I \ U], [B \ B'])$, where $x_i = z_i / \sum_{s=1}^d z_s$ for $1 \leq i \leq d$ and y is a suitable uniform distribution, whose support is contained in $\{1, \dots, d\}$ if and only if the support of z is.

Proof. Let $U = [e_{l(1)} \cdots e_{l(k)}]$, and consider the support of z with the two sets

$$S = \{i \mid z_i > 0, 1 \leq i \leq d\}, \quad T = \{j \mid z_{d+j} > 0, 1 \leq j \leq k\} \quad (4.14)$$

If S was empty, only the rows $d + j$ for $j \in T$ of player 2's payoff matrix G would be played with positive probability when player 1 uses z . However, by (4.13) each such row $[e_{l(j)}^\top \ \mathbf{0}^\top]$ has a single payoff 1 in one of the first d columns, and zeros

Chapter 4. Index and Uniqueness of Symmetric Equilibria

elsewhere, so no column of the form $d + j$ would be a best reply against z and (z, z) would not be an equilibrium of (G^\top, G) . So $S \neq \emptyset$, and x is well defined as the re-scaled vector $(z_1, \dots, z_d)^\top$.

We can assume that the unit vectors $e_{l(j)}$ for $j \in T$ are all distinct. Otherwise, if $e_{l(j)} = e_{l(j')}$ for some $j, j' \in T$ and $j \neq j'$, we could replace z by a mixed strategy \hat{z} which agrees with z except that $\hat{z}_{d+j} = z_{d+j} + z_{d+j'}$ and $\hat{z}_{d+j'} = 0$, so that j' can be omitted from T . The two mixed strategies z and \hat{z} give the same expected payoffs, $z^\top G = \hat{z}^\top G$, and give rise to the same strategy x , but \hat{z} has smaller support (this can only occur in degenerate games).

In the unit vector game $([I \ U], [B \ B'])$, the pure best replies to player 1's mixed strategy x include (and in a nondegenerate game are exactly) the following columns: $d + j$ for all $j \in T$, because the last k columns of $x^\top [B \ B']$ and of $z^\top G$ are the same except for the factor $\sum_{s=1}^d z_s$. Secondly, any column i in

$$R = S \setminus \{l(j) \mid j \in T\} \quad (4.15)$$

is a best reply to x , because for $1 \leq i \leq d$ the i th entry of $x^\top [B \ B']$ is $(z^\top G)_i / \sum_{s=1}^d z_s$ if $i \notin \{l(j) \mid j \in T\}$ (in particular, $i \in R$), or $((z^\top G)_i - z_{d+j}) / \sum_{s=1}^d z_s$ if $i = l(j)$ for some $j \in T$. In the latter case, the variable z_{d+j} is a "slack variable" for player 2's payoff in column $l(j)$, so if this column is a best reply to z , it is no longer a best reply to x .

The set of pure best replies to x therefore contains $R \cup \{d + j \mid j \in T\}$, which has the same size as the support S of x . Player 2's mixed strategy y in the unit vector game with $y_l = 1/|S|$ for $l \in R \cup \{d + j \mid j \in T\}$ and $y_l = 0$ otherwise is therefore a best reply to x . Against y , player 1, who has payoff matrix $[I \ U]$ in the unit vector game, receives payoff $1/|S|$ for each row i in R (via the i th column of I), and payoff $1/|S|$ for each row $i = l(j)$ for some $j \in T$ (via the j th column $e_{l(j)}$ of U). These are exactly the rows in the support S of x . All other rows give expected payoff zero against y . So x is a best reply to y , and (x, y) is an equilibrium of $([I \ U], [B \ B'])$, as claimed. \square

Assuming that Theorem 4.2 holds, we use the preceding lemma to prove our main result, Theorem 4.1, as follows.

Proof of Theorem 4.1. Let (B^\top, B) be a nondegenerate $d \times d$ game with symmetric

equilibrium (x, x) of symmetric index $+1$, and let P^Δ be the dual best reply polytope in (4.7) with $n = d$. By Lemma 4.7, x corresponds to a completely labelled facet F of P^Δ with opposite orientation to F_0 . By Theorem 4.2, we can add a set of vertices c_1, \dots, c_k to P^Δ , where each c_j has some label $l(j)$ in $\{1, \dots, d\}$ for $1 \leq j \leq k$, so that the labelled polytope $P_{\text{ext}}^\Delta = \text{conv}(P^\Delta \cup \{c_1, \dots, c_k\})$ has only F and F_0 as completely labelled facets. Let $U = [e_{l(1)} \cdots e_{l(k)}]$, and let the j th column of the $d \times k$ matrix B' be $c_j / (1 + \mathbf{1}^\top c_j)$, for $1 \leq j \leq k$. By Lemma 4.10, any completely labelled facet of Q , with the exception of F_0 , corresponds to a Nash equilibrium of $([I \ U], [B \ B'])$. Hence, the only such Nash equilibrium is (x, y) where $y_i = 1/|\text{supp}(x)|$ for $i \in \text{supp}(x)$ and $y_i = 0$ otherwise, which corresponds to the given symmetric equilibrium (x, x) of (B^\top, B) . Then (x, x) extended to the symmetric game (G^\top, G) with G as in (4.13) is the unique symmetric equilibrium of (G^\top, G) , because by Lemma 4.11, any other symmetric equilibrium would give rise to a different Nash equilibrium of $([I \ U], [B \ B'])$. \square

4.6 P-matrix prisms

In the remaining sections, we prove Theorem 4.2. We use a class of matrices that allows us to construct polytopes with known completely labelled facets. These are the P-matrices, which are known from mathematical programming, in particular for linear complementarity problems (Cottle et al., 1992). A $d \times d$ matrix A is a *P-matrix* if all its principal minors are positive; a principal minor of A is a determinant of the form $\det(A_{SS})$ for any subset S of $\{1, \dots, d\}$, where A_{SS} is the submatrix of A obtained by deleting all rows and columns of A that are not in S .

P-matrices are useful for our purposes because they allow the construction of a particular type of labelled polytope, the “P-matrix prism”, which has only two completely labelled facets. Here we use the notion of prism in a very general sense to denote the convex hull of two parallel simplices. Johnson et al. (2003) show that every matrix of positive index can be written as the product of three P-matrices. Given two completely labelled facets of a labelled simplicial polytope of opposite orientation, we can use this matrix decomposition to create a “stack” of three P-matrix prisms between the two facets. These prisms are placed such that each pair of intersecting prisms meets in a pair of completely labelled facets. Since each

prism only has two completely labelled facets, the only completely labelled facets of the final stack of three prisms are the two given facets of the original polytope (see Figure 4.2 in Section 4.7 for a graphic illustration of the idea).

In order for this approach to work we need to resolve two problems. First, for a general P-matrix, we cannot prove that the corresponding P-matrix prism has only two completely labelled facets. Hence we need to restrict the class of P-matrices considered. This leads to the second issue: We need to change the decomposition result in Johnson et al. (2003) to hold for this restricted class of P-matrices. In this section, we deal with these two problems (Theorem 4.12 resolves the first issue, and Proposition 4.13 the second); we restrict ourselves to P-matrices which are permutation-similar to upper triangular matrices.

For a permutation π of $\{1, \dots, d\}$, the corresponding permutation matrix is the matrix $E_\pi = [e_{\pi(1)} \cdots e_{\pi(d)}]^\top$, i.e. the matrix whose i th row is given by the $\pi(i)$ th unit vector. Multiplying a matrix $A = [a_1 \cdots a_d]$ by E_π^{-1} from the right yields $AE_\pi^{-1} = [a_{\pi(1)} \cdots a_{\pi(d)}]$, i.e. permutes the columns by π . Multiplying a matrix $A = [a_1 \cdots a_d]^\top$ from the left by E_π yields $E_\pi A = [a_{\pi(1)} \cdots a_{\pi(d)}]^\top$, which means that the rows are permuted by π . Two $d \times d$ matrices A, B are permutation-similar if there is a permutation matrix E_π such that $B = E_\pi A E_\pi^{-1}$, i.e. B is obtained from A by permuting both the rows and columns of A by π . A matrix that is permutation-similar to a P-matrix is again a P-matrix.

For a matrix C , let $\text{conv}(C)$ be the convex hull of its column vectors. Because we reserve the letters P and Q for polytopes, we denote P-matrices by letters like R, S, T .

Theorem 4.12. *Let R be a $d \times d$ matrix that is permutation-similar to an upper triangular P-matrix (i.e. an upper triangular matrix with positive diagonal entries), and that satisfies $R^\top \mathbf{1} = \lambda \mathbf{1}$ where $\lambda \neq 1$. Then the polytope $P := \text{conv}[I, R]$, where e_i and r_i have label i , only has the two “trivial” completely labelled facets $\text{conv}(I)$ and $\text{conv}(R)$.*

Proof. First, assume that $R = [r_1, \dots, r_d]$ is an upper triangular P-matrix. We will prove the claim using induction on the dimension d of the polytope P . For $d = 2$, the claim is obvious. For $d > 2$, consider the polytope generated by the points $e_1, \dots, e_{d-1}, r_1, \dots, r_{d-1}$. This polytope P' is contained in the hyperplane

$\{x \in \mathbb{R}^d \mid x_d = 0\}$, hence its dimension is less than d . By the induction hypothesis we can conclude that the only completely labelled facets of P' are $\text{conv}\{e_1, \dots, e_{d-1}\}$ and $\text{conv}\{r_1, \dots, r_{d-1}\}$.

Hence the only non-trivial completely labelled facets of $\text{conv}[I, P]$ could be the ones containing the vertices e_1, \dots, e_{d-1}, r_d or r_1, \dots, r_{d-1}, e_d . Let us first prove that there is no facet containing the first set of points. The points e_1, \dots, e_{d-1}, r_d are affinely independent, and span a hyperplane given by $\{x \mid v^\top x = 1\}$, where v is of the form $(1, \dots, 1, x)$ for

$$x = \frac{1 - r_{d1} - \dots - r_{d(d-1)}}{\lambda - r_{d1} - \dots - r_{d(d-1)}}$$

The denominator does not vanish since $\lambda - r_{d1} - \dots - r_{d(d-1)} = r_{dd} > 0$ by assumption. If $\lambda > 1$ (or < 1) then $v^\top e_d = x < 1$ (or > 1). Since $v^\top r_i = \lambda$ for $1 \leq i \leq d-1$ this means that r_1, \dots, r_{d-1} and e_d are on opposite sides of the hyperplane spanned by e_1, \dots, e_{d-1}, r_d , hence the latter set of points cannot be contained in a facet.

If the polytope $\text{conv}[I, R]$ is simplicial we are done, since then by Proposition 4.9 completely labelled facets have to come in pairs. If the polytope is not simplicial, we need to prove that there is no facet containing the affinely independent points r_1, \dots, r_{d-1}, e_d . The affine hyperplane spanned by those points is of the form $\{x \mid v^\top x = \lambda\}$ where $v = \mathbf{1} + (\lambda - 1)e_d$. Hence $v^\top e_i = 1$ for $1 \leq i \leq d-1$ and $v^\top r_d = \lambda + (\lambda - 1)r_{dd}$. Since $r_{dd} > 0$ by assumption, $v^\top r_d > \lambda$ if and only if $\lambda > 1$. This implies that r_d and e_1, \dots, e_{d-1} are on opposite sides of the hyperplane spanned by r_1, \dots, r_{d-1}, e_d , hence the latter set of points cannot be contained in a facet.

Now consider a matrix R and some permutation π so that $E_\pi R E_\pi^{-1}$ is an upper triangular P-matrix. We need to prove that the polytope P generated by the columns of R and I , where the i th columns of both matrices have label i , cannot have any completely labelled facets except for the trivial ones. This can essentially be seen using the above result for upper triangular P-matrices, since multiplication by E_π from the left is just an affine transformation, while multiplication from the right by E_π^{-1} permutes the vertices of the polytope and can be offset by a corresponding re-labelling. More precisely, assume F was a completely labelled facet of $\text{conv}[I, R]$, spanned by $e_j, j \in J$, and $r_k, k \in K$, where $J \cup K = \{1, \dots, d\}$. As made precise earlier, multiplication of a matrix by E_π^{-1} from the right is equivalent to permuting the columns of that matrix. Hence the polytope P' generated by the columns of

$R' = RE_\pi^{-1}$ and $I' = IE_\pi^{-1}$ still has F as a facet, which is spanned by the j th columns of I' for $j \in \pi^{-1}(J)$, and the k th columns of R' , for $k \in \pi^{-1}(K)$. The labelling of P induces a natural labelling on P' , where the label of the i th column of R' and I' is $\pi(i)$. Now permute the labels of the polytope P' by π^{-1} ; this is equivalent to giving the i th column of I' and R' label i . Then F is still a completely labelled facet of the relabelled polytope P' . Now we can apply the linear transformation E_π to the relabelled polytope $\text{conv}[I', R']$. Since this does not change the combinatorial structure nor the labelling of the polytope we can conclude that $E_\pi F$ is a completely labelled facet of the polytope $\text{conv}[I, E_\pi R E_\pi^{-1}]$. This facet is spanned by the unit vectors $e_j, j \in \pi^{-1}(J)$, each with label j , and by the k th columns of $E_\pi R E_\pi^{-1}$ for $k \in \pi^{-1}(K)$, each with label k . By the first part of the proof, either $\pi^{-1}(J)$ or $\pi^{-1}(K)$ must then have been empty, which in turn implies that either J or K must have been empty. \square

In Theorem 4.12, the condition that R is permutation-similar to a P-matrix is crucial: The matrix

$$R = \begin{pmatrix} 2 & 1 & 2 \\ 0 & 1 & 1 \\ 0 & 0 & -1 \end{pmatrix}$$

is upper triangular but not a P-matrix, and the polytope $\text{conv}[I, R]$ has four completely labelled facets.

The following useful result is due to Johnson et al. (2003), who proved that every matrix with positive determinant is the product of at most three P-matrices. Using the same proof, their result can easily be modified so that it applies to P-matrices that are permutation-similar to upper triangular matrices, which we need for our construction.

Proposition 4.13. *Every non-diagonal matrix A with positive determinant is the product of exactly three matrices*

$$A = RST$$

where R, S and T are permutation-similar to upper triangular P-matrices.

We essentially follow the proof of Johnson et al. (2003, Theorem 2.6) for non-diagonal matrices, and point out where we keep track of the shape of the P-matrices.

We first need the following lemma, analogous to Johnson et al. (2003, Lemma 2.4). A $d \times d$ matrix has a *nested sequence of positive principal minors* if it contains a sequence of positive principal minors of descending order $d, \dots, 1$, such that each minor's index set contains the next.

Lemma 4.14. *For every non-diagonal $d \times d$ matrix A with positive determinant there exists a matrix R such that AR has a nested sequence of positive principal minors, and R is permutation-similar to a lower triangular P-matrix.*

Proof. The proof is by induction on d . Since A is non-diagonal, we can find a permutation matrix E_π such that $A' = E_\pi A E_\pi^{-1}$ has its entry a'_{12} non-zero. It now suffices to find a matrix R , permutation-similar to an upper triangular P-matrix, such that $A'R$ has a nested sequence of positive principal minors, as then does $E_\pi^{-1}(A'R)E_\pi = A(E_\pi^{-1}RE_\pi)$.

For $d = 2$, choose r such that $a'_{11} + a'_{12}r > 0$. Then for

$$R = \begin{pmatrix} 1 & 0 \\ r & 1 \end{pmatrix}$$

$A'R$ has a nested sequence of positive principal minors.

Now for $d > 2$, let

$$\hat{R} = \begin{pmatrix} 1 & 0 \\ r & I \end{pmatrix}$$

where I is the $(d-1) \times (d-1)$ identity matrix, and $r = (r_1, r_2, 1, \dots, 1)$ for $r_1, r_2 \in \mathbb{R}$ that we will have to choose suitably. We get

$$\hat{R}^{-1} = \begin{pmatrix} 1 & 0 \\ -r & I \end{pmatrix}$$

Now partition

$$\hat{R}^{-1}A'^{-1} = \begin{pmatrix} b & v^\top \\ u & B \end{pmatrix}$$

where B is a $(d-1) \times (d-1)$ -matrix. Write $A'^{-1} = (\alpha_{ij})_{1 \leq i, j \leq d}$. If $\alpha_{12} = 0$, the second column of A'^{-1} must have one non-zero off-diagonal entry: The product of the first line of A' with the second column of A'^{-1} must be zero, which implies that if all off-diagonal entries of the second column of A'^{-1} were zero, the whole

column would have to be zero since $a'_{12} \neq 0$. Hence for the first column of B we get that its entry $b_{i1} = \alpha'_{(i+1)2} \neq 0$ for some $i \geq 2$. If $\alpha_{12} \neq 0$, we can choose r_2 such that

$$b_{21} = -r_2\alpha_{12} + \alpha_{32} \neq 0$$

Hence in both cases B is non-diagonal. The (1,1) entry of $A'\hat{R}$ is $a'_{11} + a'_{12}r_1 + a'_{13}r_2 + a'_{14} + \dots + a'_{1d}$. Since $a'_{12} \neq 0$, we can choose r_1 such that this sum is positive. Since $\det(\hat{R}^{-1}A'^{-1}) > 0$, Cramer's rule implies that $\det(B) > 0$. Hence $\det(B^{-1}) > 0$, and B^{-1} is also non-diagonal.

By induction hypothesis there is a $(d-1) \times (d-1)$ matrix S , permutation-similar to a lower triangular P-matrix, such that $B^{-1}S$ has a nested sequence of positive principal minors. Hence so has $S^{-1}B$, since any principal minor of a non-singular matrix M equals the determinant of M , multiplied by the complementary principal minor of the inverse matrix M^{-1} (Cottle et al., 1992). Moreover, S^{-1} is permutation-similar to a lower triangular P-matrix (in particular, it has positive determinant). We get

$$\begin{pmatrix} 1 & 0 \\ 0 & S^{-1} \end{pmatrix} \begin{pmatrix} 1 & 0 \\ -r & I \end{pmatrix} A'^{-1} = \begin{pmatrix} 1 & 0 \\ 0 & S^{-1} \end{pmatrix} \begin{pmatrix} b & v^\top \\ u & B \end{pmatrix} = \begin{pmatrix} b & v^\top \\ S^{-1}u & S^{-1}B \end{pmatrix}$$

where the latter matrix has positive determinant, hence by choice of S a nested sequence of positive principal minors. But the product of the first two matrices is

$$R^{-1} = \begin{pmatrix} 1 & 0 \\ -S^{-1}w & S^{-1} \end{pmatrix}$$

which is permutation-similar to a lower triangular P-matrix since S^{-1} is. □

Proof of Proposition 4.13. By Lemma 4.14, there exists a matrix T such that AT^{-1} has a nested sequence of positive principal minors, and T^{-1} is permutation-similar to a lower triangular P-matrix (hence so is T). This means that there is a permutation matrix E_π such that $E_\pi AT^{-1} E_\pi^{-1}$ has a leading sequence of positive principal minors. Then $E_\pi AT^{-1} E_\pi^{-1}$ has a LU -factorization $E_\pi AT^{-1} E_\pi^{-1} = LU$, where L and U are lower and upper triangular P-matrices, respectively (Cottle et al., 1992). We get

$$A = E_\pi^{-1} L U E_\pi T = (E_\pi^{-1} L E_\pi) (E_\pi^{-1} U E_\pi) T$$

Since every lower triangular P-matrix is permutation-similar to an upper triangular P-matrix, the result follows for $R = E_\pi^{-1} L E_\pi$ and $S = E_\pi^{-1} U E_\pi$. \square

4.7 Re-arranging the polytope P^Δ

In the following sections, we give a proof of our main theorem in the polytope version, Theorem 4.2. We want to use the P-matrix prisms from the previous section, whose completely labelled facets we know. Our goal is to create a stack of such prisms between the two completely labelled facets F and F_0 of the polytope P^Δ in Theorem 4.2. One of the completely labelled facets of a P-matrix prism is always the facet $\text{conv}(I)$, which is spanned by the unit vectors. To adapt to this, we have to move the polytope P^Δ such that F_0 is spanned by the unit vectors as well. For this reason we formulate a slightly different version of Theorem 4.2 in Proposition 4.15, where P^Δ is transformed accordingly. After this transformation, the facet F has positive orientation (this will become clear later when we explain the transformation in more detail).

The idea of the proof then is as follows: If we write $F = \text{conv}(C)$, where the columns of C are ordered according to their labelling, C must have positive determinant. Using Proposition 4.13, we write $C = RST$ as a product of three matrices that are permutation similar to upper triangular P-matrices. Using these matrices, we generate a stack of three polytopes between the facets F and F_0 , such that each polytope in the stack has only two completely labelled facets, i.e. its top and bottom. More precisely, we use the three polytopes $\text{conv}[I, R]$, $\text{conv}[R, RS]$ and $\text{conv}[RS, C]$ for our stack. All that remains to do is to expand the facets $\text{conv}(R)$ and $\text{conv}(RS)$ of the polytopes in the stack to “catch” all of the polytope P^Δ in the interior of the extended polytope, except for the top facet F and bottom facet F_0 . For a visualisation see Figure 4.2.

One of the crucial points for this idea to work is that the facets F and F_0 have to be parallel. This is only possible if the two facets are disjoint; if they share certain points, we have to impose an analogous technical condition, stated as equation (4.16) in Proposition 4.15.

In the remainder of this thesis, we will often have to refer to the columns of a given

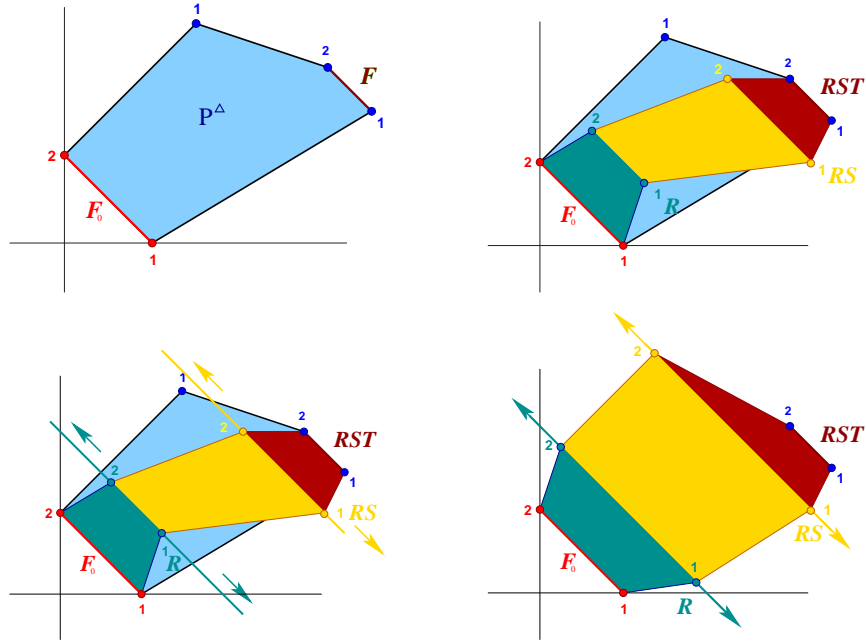


Figure 4.2: Example of using Theorem 4.12 and Proposition 4.13 to create a stack of polytopes that by extension makes the facets F and F_0 the only completely labelled facets of the extended polytope.

matrix. Unless stated otherwise, we use the corresponding lowercase letter for the columns of a matrix, i.e. c_i for a matrix C etc.

Proposition 4.15. *Let $P^\Delta = \text{conv}[I, C, C']$ be a labelled simplicial polytope, where C and C' are positive $d \times k$ and $d \times n$ matrices respectively, for some $2 \leq k \leq d$ and $n \geq 0$. Assume that all columns of C and C' are contained in the open half-space $\{x \mid \mathbf{1}^\top x > 1\}$, and both e_i (for $1 \leq i \leq d$) and c_i (for $1 \leq i \leq k$) have label i (we do not need any condition on the labels of C'). Suppose we are given a positively oriented completely labelled facet*

$$F = \text{conv}\{c_1, \dots, c_k, e_{k+1}, \dots, e_d\}$$

Denote by $\begin{pmatrix} 0 \\ \mathbf{1} \end{pmatrix}$ the vector with 0's in the first k coordinates and 1's in the others, and by $\begin{pmatrix} \mathbf{1} \\ 0 \end{pmatrix}$ the vector $\mathbf{1} - \begin{pmatrix} 0 \\ \mathbf{1} \end{pmatrix}$. Assume that F can be written as

$$F = \{x \in P^\Delta \mid (\alpha \begin{pmatrix} \mathbf{1} \\ 0 \end{pmatrix} + \begin{pmatrix} 0 \\ \mathbf{1} \end{pmatrix})^\top x = 1\} \quad (4.16)$$

for some $\alpha < 1$, where $(\alpha \begin{pmatrix} \mathbf{1} \\ 0 \end{pmatrix} + \begin{pmatrix} 0 \\ \mathbf{1} \end{pmatrix})^\top x \leq 1$ for $x \in P^\Delta$. Then we can add labelled vertices to P^Δ such that the only completely labelled facets of the extended polytope

Chapter 4. Index and Uniqueness of Symmetric Equilibria

are F and $F_0 = \text{conv}(I)$. It suffices to add $k(d - k + 1)$ vertices if $k < d$, and $2d$ vertices if $k = d$.

It is easy to see that Proposition 4.15 implies Theorem 4.2. Before we see how this works, two useful results:

Lemma 4.16. *For $\delta > 1/d$ and E the matrix having all entries equal to 1, the affine transformation*

$$\mathbb{R}^d \rightarrow \mathbb{R}^d, \quad x \mapsto (\delta E - I)^{-1}x$$

is orientation preserving if d is odd, and orientation reversing otherwise.

Proof. It suffices to calculate the determinant of $\delta E - I$. By multi-linearity of the determinant and Laplace expansion, this determinant is easily seen to be $\delta d(-1)^{d-1} + (-1)^d$, which is positive (negative) if d is odd (even). \square

Lemma 4.17. *Every pure strategy equilibrium in a non-degenerate bimatrix game can be made the unique equilibrium by adding one strategy for the column player. The payoff column for the row player can be chosen to be a suitable unit vector.*

Proof. This is straightforward; see von Schemde (2005, Lemma 4.1). \square

Now we prove that Proposition 4.15 implies Theorem 4.2

Proof of Theorem 4.2. Assume we are given a labelled simplicial d -polytope P^Δ with $\mathbf{0}$ in its interior, and let F_0 and F be two completely labelled facets of opposite orientation. By linearly transforming the polytope P^Δ , we can assume without loss of generality that the completely labelled facet F_0 is spanned by the negative unit vectors, each labelled canonically, while $\mathbf{0}$ is still contained in the polytope. By a coordinate change and a relabelling of vertices we can assume the completely labelled facet F to be of the form $F = \text{conv}\{c_1, \dots, c_k, -e_{k+1}, \dots, -e_d\}$ for some $1 \leq k \leq d$, where c_i has label i , and the negative unit vectors still have their canonical labelling. Let $C = [c_1 \cdots c_k]$, and denote by C' the (potentially empty) matrix of the remaining vertices of P^Δ , i.e. the vertices neither in F nor F_0 . We are now in the situation of a polytope given by a unit vector game as in Lemma 4.10. Hence the completely labelled facets of P^Δ correspond to the equilibria of

a bimatrix game $([U \ U'], [B \ B'])$, where B, B' are rescaled versions of C, C' , and U, U' consist of unit vectors.

For the case $k = 1$, there is a simple game-theoretic proof for the fact that we can make the two facets F and F_0 the only completely labelled facets of the polytope by adding just one strategy (see Lemma 4.17). Hence for the remainder of the proof, we can assume $k \geq 2$.

We can rescale the rows of $[B \ B']$, each by a positive scalar, such that the strategy of the row player in the equilibrium corresponding to the facet F is uniformly distributed. This corresponds to multiplying $[B \ B']$ from the left by a suitable diagonal matrix with positive diagonal entries. This procedure does not change the combinatorial structure of P , hence neither that of P^Δ : In the definition of P , we can easily replace B by a multiple DB for some positive diagonal matrix D , and obtain a linearly equivalent polytope. Also, by a suitable choice of rescaling, we can assume that $\mathbf{0}$ is still contained in P^Δ .

After rescaling, the facet F is given by $\{x \in P^\Delta \mid (\bar{\alpha} \binom{1}{\mathbf{0}} - \binom{\mathbf{0}}{1})^\top x = 1\}$ for some $\bar{\alpha} > -1$, with $(\bar{\alpha} \binom{1}{\mathbf{0}} - \binom{\mathbf{0}}{1})^\top x < 1$ for all vertices x of P^Δ that are not in F . By applying the affine transformation $T_\mu : x \mapsto (I + \mu E)x + \mu \mathbf{1} = x + \mu(\mathbf{1}^\top x + 1)\mathbf{1}$, for μ big enough, we can assume without loss of generality that $C, C' > 0$. This affine transformation leaves F_0 invariant and stretches the rest of the polytope towards infinity, while keeping 0 in the interior of the polytope. Since 0 remains in the polytope, during the transformation no facet-defining hyperplane crosses the origin, hence F and F_0 keep their orientation.

Note that a linear transformation of a polytope by a non-singular matrix M changes a normal v on a facet F to $(M^\top)^{-1}v$, which is the normal on MF . Translating a polytope changes the normal on a facet by a positive real scalar as long as the facet-defining hyperplane does not cross the origin. Hence in our case, the normal on the transformed facet $T_\mu(F)$ is again of the form $\bar{\alpha} \binom{1}{\mathbf{0}} - \binom{\mathbf{0}}{1}$, for some new $\bar{\alpha} > -1$.

We need to put the polytope P^Δ “in the right position” by moving F_0 to $\text{conv}(I)$, and F to some facet of positive orientation. To achieve this, we cannot use the “obvious” linear transformation $-I$, since the transformed facet $-IF$ would have negative orientation. Instead, we choose the following affine transformation: We add $(\frac{1}{d} + \varepsilon)\mathbf{1}$ to all vertices for some small $\varepsilon > 0$, and then apply the linear transforma-

tion that maps $\mathbf{1}(\frac{1}{d} + \varepsilon) - e_i$ to e_i . By abuse of notation, we denote the transformed facets again by F and F_0 , respectively. We can choose ε small enough such that F does not change orientation during the translation. This means that the facet defining hyperplane which defines F does not cross the origin, which also implies that, by similar considerations as above, this hyperplane is transformed by the translation to $\{x \mid (\bar{\alpha} \binom{1}{0} - \binom{0}{1})^\top x = \lambda\}$ for some $\lambda > 0$, and by the subsequent linear transformation to $\{x \mid (\alpha \binom{1}{0} - \beta \binom{0}{1})^\top x = \lambda\}$ for some α, β , with strict inequality for vertices not in F . We might even assume $\lambda = 1$, which implies $\beta = 1$ (since $e_d \in F$ if $k < d$), and $\alpha < 1$ since $\alpha = \left(\alpha \binom{1}{0} + \binom{0}{1}\right) e_1 < 1$ (using that $e_1 \notin F$).

Finally, after applying the linear transformation $((\frac{1}{d} + \varepsilon)E - I)^{-1}$ to the translated polytope, by Lemma 4.16 both F and F_0 have positive orientation. Hence we can apply Proposition 4.15 and add labelled vertices to the transformed polytope to make the facets F_0 and F the only completely labelled facets. Reversing all transformations does not change the combinatorial structure of the polytope. This proves that Proposition 4.15 implies Theorem 4.2. \square

4.8 Disjoint completely labelled facets

All that remains to be done is proving Proposition 4.15. This proof takes up the remainder of this chapter. As explained in the previous section, we would like to insert a stack of P-matrix prisms between the two completely labelled facets F and F_0 . For this, we need the two completely labelled facets of P^Δ to be parallel (Figure 4.2 in Section 4.7 provides an intuition for the reasons behind this). This can only be achieved if the two facets are disjoint. For this reason, we first give a proof for the case of disjoint completely labelled facets in the present section. In terms of symmetric equilibria, disjoint facets correspond to the case of a symmetric equilibrium of full support. The general case is treated in Section 4.9.

We now prove the following result, which is slightly stronger than Proposition 4.15 for disjoint completely labelled facets.

Proposition 4.18. *For $d \geq 2$, consider a labelled d -polytope $P^\Delta = \text{conv}[I, C, C']$ where $C \in \mathbb{R}^{d \times d}$ and $C' \in \mathbb{R}^{d \times n}$ (for some $n \geq 0$). Assume there is some $\lambda > 1$ such that $C^\top \mathbf{1} = \lambda \mathbf{1}$ and $\mathbf{1} < C'^\top \mathbf{1} < \lambda \mathbf{1}$. Then P^Δ has two parallel disjoint completely*

labelled facets $F_0 = \text{conv}(I)$ and $F = \text{conv}(C)$. If F is positively oriented, we can add a set of $2d$ pairwise distinct labelled points $X = \{x_1, \dots, x_{2d}\}$ to the polytope P^Δ such that the following conditions hold:

- (a) $P_{\text{ext}}^\Delta = \text{conv}[I, C, C', X]$ has only two completely labelled facets, $\text{conv}(I)$ and $\text{conv}(C)$ (where by abuse of notation we write X for the matrix $[x_1 \dots, x_{2d}]$).
- (b) Any column of C' is caught in the relative interior of the convex hull of the added points, i.e. $C' \subset \text{relint}(\text{conv}(X))$.
- (c) Given $\varepsilon_1, \varepsilon_2 > 0$ small enough, we can choose the set X such that it consists of two subsets X_1 and X_2 , each of cardinality d , such that for all $x \in X_1$, $\mathbf{1}^\top x = 1 + \varepsilon_1$ and for all $x \in X_2$, $\mathbf{1}^\top x = \lambda - \varepsilon_2$, and P_{ext}^Δ is the union of the three polytopes $\text{conv}(I, X_1)$, $\text{conv}(X_1, X_2)$ and $\text{conv}(X_2, C)$.
- (d) All points in X are extremal points of P_{ext}^Δ .

Note that this result is slightly stronger than what we actually need for the purposes of this section: In order to prove Proposition 4.15, we could omit conditions (b)-(d) and assume that $C, C' > 0$. However, we will need this stronger version to extend the proof to polytopes with non-disjoint completely labelled facets in Section 4.9.

Before we prove the proposition, we collect a few ingredients for the proof. The following Lemma is needed for the top and bottom facets of each of the “stack” polytopes to be parallel. ¹

Lemma 4.19. *If a matrix M is permutation-similar to an upper triangular P-matrix, then so are DM and MD for any diagonal matrix D with positive entries.*

Proof. The claim is obvious for M an upper triangular P-matrix. For M permutation-similar to such a matrix, the claim follows directly using the following observation: For any permutation matrix E_π and any diagonal matrix D we get

$$DE_\pi = E_\pi D' \tag{4.17}$$

for a suitable diagonal matrix D' . This is due to the fact that multiplying E_π by D from the left results in the rows of E_π being scaled by the respective diagonal

¹Whenever we use the term “stack(ed) polytope” in this chapter, we do not refer to the different technical term as used in polytope theory (see Ziegler, 1995), but mean it in our illustrative sense, hoping that this does not lead to confusion.

entries of D , whereas multiplying E_π with D' from the right means scaling the columns of E_π . Using that E_π has only one nonzero entry in each row and column, (4.17) follows easily. \square

The following three results will be used for expanding the facets $\text{conv}(R)$ and $\text{conv}(RS)$.

Proposition 4.20. *Consider a $d \times d$ matrix M such that $M^\top \mathbf{1} = \lambda \mathbf{1}$ for some constant $\lambda \neq 1$. For any vector s such that $s^\top \mathbf{1} = 0$, the shear that leaves the unit vectors invariant and maps m_i to $m_i + s$ is a bijective affine transformation, hence leaves the combinatorial structure of the polytope $\text{conv}[I, M]$ invariant.*

Proof. Define an affine transformation $T : \mathbb{R}^d \rightarrow \mathbb{R}^d, x \mapsto x + \frac{1}{d} \mathbf{1}$. Consider the line $L = \mathbb{R} \mathbf{1}$ and its orthogonal complement, $L^\perp = \{x \mid x^\top \mathbf{1} = 0\}$. Let $f : L \rightarrow L^\perp$ be the linear map given by $\alpha \mathbf{1} \mapsto s \frac{d\alpha}{\lambda - 1}$. Then $S : L^\perp \oplus L \rightarrow L^\perp \oplus L, (x, y) \mapsto (x + f(y), y)$ is a bijective linear map (a linear shear). Now the desired shear, which leaves the unit vectors invariant and moves m_i to $m_i + s$, is given by $T \circ S \circ T^{-1}$, and the claim follows. \square

In order to state the next results we need to remind the reader of the concept of a projective map. For an introduction to projective maps in the context of polytopes consult Grünbaum (2003) or Ziegler (1995), of which we give a very short summary here. A projective map τ on \mathbb{R}^d , given by a $d \times d$ matrix Z , vectors $a, z \in \mathbb{R}^d$, and some real number a_{d+1} , is defined as

$$\tau : \{x \in \mathbb{R}^d \mid a^\top x + a_{d+1} \neq 0\} \rightarrow \mathbb{R}^d, x \mapsto \frac{Zx + z}{a^\top x + a_{d+1}} \quad (4.18)$$

If the matrix

$$\begin{pmatrix} Z & z \\ a^\top & a_{d+1} \end{pmatrix} \quad (4.19)$$

is non-singular, such a projective map is called a projective transformation. If $a = \mathbf{0}$, such a projective transformation reduces to an affine transformation. We say that a projective transformation τ is valid for a polytope $P \subset \mathbb{R}^d$ if P is contained in one of the two half-spaces on which τ is defined. Then, $\tau(P)$ is combinatorially equivalent to P . This view of projective maps suffices for our purposes, but a projective map as defined in (4.18) can also be understood as the map of projective space $\mathbb{P}^d \mathbb{R} \rightarrow \mathbb{P}^d \mathbb{R}$ arising from the linear map on \mathbb{R}^{d+1} given in (4.19). For a

short and accessible treatment of projective space and projective maps see Gallier (2001).

Proposition 4.21. *Consider a matrix M such that $M^\top \mathbf{1} = \lambda \mathbf{1}$ for some constant $\lambda > 1$. For any real numbers $t > 1$ and $u \neq \frac{1-t\lambda}{d}$, there is a projective transformation that is valid for the polytope $\text{conv}(I, M)$, and maps m_i to $tm_i + u\mathbf{1}$ while leaving the unit vectors invariant.*

Proof. Choose $\varepsilon = \frac{\lambda-1}{t-1}$. The projective transformation given by

$$x \mapsto \frac{(1-\lambda-\varepsilon)x}{\mathbf{1}^\top x - \lambda - \varepsilon}$$

is defined on the polytope $\text{conv}[I, M]$ for $\varepsilon > 0$, i.e. $t > 1$. It maps each unit vector to itself, and each vector m_i to tm_i .

Now define $\mu = \frac{u}{t\lambda-1}$ (since $t, \lambda > 1$, the denominator does not vanish). Define an affine map

$$x \mapsto (I + \mu E)x - \mu \mathbf{1}$$

which is bijective if and only if $\mu \neq -1/d$, or $u \neq \frac{1-t\lambda}{d}$. This transformation leaves the unit vectors invariant, and maps tm_i to $tm_i + u\mathbf{1}$. The concatenation of the two maps yields the desired projective transformation. \square

Corollary 4.22. *Consider non-singular matrices M, N such that $M^\top \mathbf{1} = \lambda \mathbf{1}$, $N^\top \mathbf{1} = \lambda' \mathbf{1}$ for some constants $\lambda \neq \lambda'$. Choose any point s in the hyperplane $\{x \mid x^\top \mathbf{1} = \lambda\}$. Then for any $t > 0$ and $M' = M + t(M - [s \cdots s])$, the polytope $\text{conv}[M', N]$ is projectively equivalent to $\text{conv}[M, N]$. This means that we can replace the vertices m_i by $m_i + t(m_i - s)$ without changing the combinatorial structure of the polytope $\text{conv}[M, N]$. Geometrically, this corresponds to “blowing up” the facet $\text{conv}(M)$ from the point of reference s .*

Proof. By applying the linear transformation N^{-1} to the polytope $\text{conv}[M, N]$, we can assume without loss of generality that $N = I$ and $\lambda \neq 1$. We can even assume $\lambda > 1$, since otherwise we can reflect the polytope in the hyperplane $\{x \mid x^\top \mathbf{1} = 1\}$, using the reflection

$$T : x \mapsto x - \frac{2}{d}(\mathbf{1}^\top x - 1) \cdot \mathbf{1}$$

and applying our result to the polytope $\text{conv}[TM, I]$, with the reflected reference point Ts .

For the case where $N = I$ and $\lambda > 1$, the idea is as follows: First we apply a shear that moves s to $\frac{\lambda}{d}\mathbf{1}$, then we blow up the sheared polytope using the projective transformation in Proposition 4.21, after which we need to undo the shear. More precisely, define $s' = \frac{\lambda}{d}\mathbf{1} - s$, and consider the shear that leaves the unit vectors invariant and maps m_i to $m_i + s'$. Define $t' = t + 1$ and $u = (1 - t')\lambda/d$. Then $t' > 1$ and $u \neq \frac{1-t'\lambda}{d}$, and the projective transformation in Proposition 4.21 maps $m_i + s'$ to $t'(m_i + s') + u\mathbf{1}$. Undoing the first shear maps $t'(m_i + s') + u\mathbf{1}$ to $t'(m_i + s') + u\mathbf{1} - s'$. The concatenation of these three maps leaves the unit vectors invariant, and maps m_i to

$$\begin{aligned} t'(m_i + s') + u\mathbf{1} - s' &= t'(m_i + \frac{\lambda}{d}\mathbf{1} - s) + (1 - t')\frac{\lambda}{d}\mathbf{1} - \frac{\lambda}{d}\mathbf{1} + s \\ &= t'm_i - t's + s \\ &= m_i + tm_i - ts \end{aligned}$$

Hence we get the desired map as the concatenation of a projective transformation with two affine isomorphisms, and the claim follows. \square

Proof of Proposition 4.18. To provide geometric intuition we have sketched the proof in Figure 4.2. We assume the columns of C to be given in the order of their labels. This implies that $\det(C) > 0$ since F has positive orientation. If C is non-diagonal, we can by Proposition 4.13 write C as the product of exactly three matrices $C = RST$, which are each permutation-similar to an upper triangular P-matrix. If C is diagonal, it must be of the form λI , hence can obviously be written as such a product as well. Without loss of generality we can assume that the columns of the P-matrices are scaled such that they add up to a suitable positive constant. This can be achieved by choosing suitable positive diagonal matrices D, D' and writing

$$C = (RD)(D^{-1}SD')(D'^{-1}T) \tag{4.20}$$

which by Lemma 4.19 does not change the fact that the factors in this product are permutation-similar to upper triangular P-matrices.

Let $R' = RS$. For any choice of λ_1, λ_2 with $1 < \lambda_1 < \lambda_2 < \lambda$, we can assume by (4.20) that $R'^\top \mathbf{1} = \mathbf{1}\lambda_1$ and $R'^\top \mathbf{1} = \mathbf{1}\lambda_2$. We can choose λ_1 and λ_2 such that all vertices C' which are not contained in the completely labelled facets are contained in the set

$$\{y \in \mathbb{R}^d \mid \lambda_1 < y^\top \mathbf{1} < \lambda_2\} \tag{4.21}$$

By Theorem 4.12, each of the three “stack” polytopes $\text{conv}[I, R]$, $\text{conv}[R, R']$ and $\text{conv}[R', C]$ has only two completely labelled facets, that is its “top” and “bottom” facet. Essentially, all we need to do now is blow up the two middle facets of the stack, i.e. $\text{conv}(R)$ and $\text{conv}(R')$, so that their convex hull (together with I and C) contains the polytope P^Δ , and then make sure that we get a proper “stack” of polytopes, i.e. that the convex hull of I, R, R' and C is indeed the union of the three stack polytopes.

By Corollary 4.22, we can blow up the facet $\text{conv}(R)$ from its barycenter s , i.e. replace it by

$$R_t = (1+t)R - t[s \dots s]$$

for some $t \geq 0$, without changing the combinatorial structure of any of the stack polytopes. We can then “translate” the blown-up facet $\text{conv}(R_t)$ into different hyperplanes by adding a suitable scalar multiple of $\mathbf{1}$ to the facet. This does not change the combinatorial structure of any of the stack polytopes (as long as we do not put any pair of facets into the same hyperplane, which would “squash” the corresponding stack polytope). This is due to the fact that translating one of the facets of a stack polytope by a multiple of $\mathbf{1}$ corresponds to applying the affine map

$$x \mapsto (I + \mu E)x - \mu v \mathbf{1} = x + \mu(\mathbf{1}^\top x - v)\mathbf{1} \quad (4.22)$$

for suitable choices of μ and v . Denote the translation of R_t into the hyperplane $\{x \mid x^\top \mathbf{1} = \lambda'\}$ by $R_{t, \lambda'}$. More precisely,

$$R_{t, \lambda'} = R_t + \frac{\lambda' - \lambda_1}{d} \mathbf{1}$$

Choose t big enough such that for every λ' in the closed interval $[1, \lambda]$, the convex hull of I, C and $R_{t, \lambda'}$ contains P^Δ . Such a t must exist, since the function

$$[1, \lambda] \rightarrow \mathbb{R}, \lambda' \mapsto \inf\{t \in [0, \infty) \mid P^\Delta \subset \text{conv}[I, C, R_{t, \lambda'}]\}$$

is continuous, hence bounded. By slightly increasing t , we can assume that not only P^Δ is contained in $\text{conv}[I, C, R_{t, \lambda'}]$ for every $\lambda' \in [1, \lambda]$, but that the vertices C' of P^Δ that are neither on F nor on F_0 are even contained in the relative interior of that convex hull. Moreover, we can assume t to be sufficiently large such that F_0 is contained in the relative interior of $\text{conv}(R_{t, 1})$, the translation of $\text{conv}(R_t)$

into the hyperplane $\{x \mid x^\top \mathbf{1} = 1\}$. Finally, we can assume that the analogous statement holds for the facets R' and F : Denote by R'_t the blow-up of R' from the barycenter of $\text{conv}(R')$ by the factor t , and by $R'_{t,\lambda}$ the translation of R'_t to the hyperplane $\{x \mid x^\top \mathbf{1} = \lambda\}$. Then we can assume t to be sufficiently large such that F is contained in the relative interior of $\text{conv}(R'_{t,\lambda})$.

Our choice of the expansion factor t implies that P^Δ is contained in the polytope $\text{conv}[I, C, R_t, R'_t]$, and C' in its relative interior. However, so far our construction does not guarantee that the polytope $\text{conv}[I, C, R_t, R'_t]$ is the union of the three stack polytopes $\text{conv}[I, R_t]$, $\text{conv}[R_t, R'_t]$ and $\text{conv}[R'_t, C]$, whose facet structure we know (in terms of completely labelled facets). To achieve this, we need to move the middle facets $\text{conv}(R_t)$ and $\text{conv}(R'_t)$ sufficiently outwards. By our choice of expansion parameter t we can translate the facet $\text{conv}(R_t)$ arbitrarily close to the facet F_0 , and the facet $\text{conv}(R'_t)$ arbitrarily close to the facet F , while keeping the polytope P^Δ inside the corresponding convex hull (and C' in its relative interior), and without changing the combinatorial structure of any of the stack polytopes. By replacing R_t by $R_{t,1+\varepsilon_1}$, and R'_t by $R'_{t,\lambda-\varepsilon_2}$ for some suitably small $\varepsilon_1, \varepsilon_2 > 0$, we ensure that when we put the three stacks polytopes together, they do not start “interfering with each other”: No point in F_0 can see any point in the upper stack polytope $\text{conv}[R'_{t,\lambda-\varepsilon_2}, C]$, nor does any vertex in F see a point in the lower stack polytope $\text{conv}[I, R_{t,1+\varepsilon_1}]$ (where we say that two points in $\text{conv}[I, C, R_{t,1+\varepsilon_1}, R'_{t,\lambda-\varepsilon_2}]$ “see” each other if the convex hull of those two points does not intersect the relative interior of the polytope).

We can conclude that for any small enough choice of $\varepsilon_1, \varepsilon_2 > 0$, we get that

$$\begin{aligned} Q &= \text{conv}[I, C, R_{t,1+\varepsilon_1}, R'_{t,\lambda-\varepsilon_2}] \\ &= \text{conv}[I, R_{t,1+\varepsilon_1}] \cup \text{conv}[R_{t,1+\varepsilon_1}, R'_{t,\lambda-\varepsilon_2}] \cup \text{conv}[R'_{t,\lambda-\varepsilon_2}, C] \end{aligned}$$

i.e. the first polytope is indeed the union of three stack polytopes, as desired. We can conclude that the facets of Q are given by the facets of the three stack polytopes, apart from those that have been glued together (which are the two middle facets of the stack). Since by our choice of t , Q contains every vertex C' in its relative interior, this implies that for $X = [R_{t,1+\varepsilon_1}, R'_{t,\lambda-\varepsilon_2}]$, the extended polytope $P_{\text{ext}}^\Delta = \text{conv}[I, C, C', X]$ does not have any completely labelled facets apart from F and F_0 , which proves statement (a) of the Proposition. Statements (b)-(d) are obvious from our construction. \square

4.9 General completely labelled facets

In this section we prove Proposition 4.15 for the case in which the two completely labelled facets are not disjoint. This proof concludes our proof of Theorem 4.2.

The main challenge in this case is that we cannot stack P-matrix prisms between the two facets F and F_0 of P^Δ , since they share vertices. To overcome this problem, we project the “disjoint” part of the polytope into a lower dimensional space. The projected polytope inherits the labelling from P^Δ , and has two completely labelled facets, which are the projections of F and F_0 . These two projected facets are disjoint and parallel. This means that we can apply Proposition 4.18 to add points to this lower-dimensional polytope to make the projected facets unique.

We then need to “lift” the added vertices back into the higher dimensional space, so that we can use them as added vertices for the original polytope P^Δ . For technical reasons, this process of “lifting” creates several copies of each vertex added in the lower dimension. Hence the number of vertices added to P^Δ is no longer bounded by a linear function, as in the previous section, but grows quadratically.

We use the following projective projection to create a lower-dimensional polytope from P^Δ :

$$p : \{x \in \mathbb{R}^d \mid \begin{pmatrix} \mathbf{0} \\ \mathbf{1} \end{pmatrix}^\top x \neq 1\} \rightarrow \mathbb{R}^d, x \mapsto \frac{I_k x}{1 - \begin{pmatrix} \mathbf{0} \\ \mathbf{1} \end{pmatrix}^\top x} \quad (4.23)$$

where I_k is the $d \times d$ matrix $[e_1 \cdots e_k \ \mathbf{0} \cdots \mathbf{0}]$. We would like to apply this projection to the polytope $\text{conv}[I, C, C']$ in Proposition 4.15 to get a lower-dimensional polytope with disjoint completely labelled facets. The projection is not defined on the last $d - k$ unit vectors (which are the shared vertices of F and F_0), hence we need to restrict it to the remaining vertices of P^Δ . Figure 4.3 illustrates how the projection transforms a d -polytope into a lower dimensional polytope.

In the following lemma, we analyze the facet structure of the projected polytope.

Lemma 4.23. *Consider a d -polytope $Q = \text{conv}[I, A]$ for some $d \times n$ -matrix $A > 0$ that satisfies $\begin{pmatrix} \mathbf{0} \\ \mathbf{1} \end{pmatrix}^\top A < \mathbf{1}$, meaning that the projection p in (4.23) is defined on all columns of A . Define Q_k to be the k -polytope given as the convex hull of the points $e_1, \dots, e_k, p(a_1), \dots, p(a_n)$. Then every facet F of Q that consists of the last $d - k$ unit vectors e_{k+1}, \dots, e_d and some other vertices x_1, \dots, x_s (note that Q need not be simplicial) yields a facet $F_k = \text{conv}\{p(x_1), \dots, p(x_s)\}$ of the polytope Q_k .*




Figure 4.3: Example of a projection as in (4.23) applied to a labelled simplicial polytope of the form $\text{conv}[I, C, C']$ as in Proposition 4.15, where $d = 3$ and $k = 2$. The 3-polytope has several completely labelled facets, among these the unit simplex, which is the large dashed triangle on the “back” of the polytope, and the shaded facet. The shaded 2-polytope is the projection of this 3-polytope; the thin dashed lines indicate the projection lines.

Proof. First, observe that Q_k is indeed a k -dimensional polytope. This is due to the fact that for any set of points $x_1, \dots, x_r \in \{a_1, \dots, a_n, e_1, \dots, e_k\}$ any affine dependence of the $p(x_i)$, $\sum_{i=1}^r \gamma_i p(x_i) = 0$ with $\sum_{i=1}^r \gamma_i = 0$, gives rise to an affine dependence of $x_1, \dots, x_r, e_{k+1}, \dots, e_d$ via

$$0 = \sum_{i=1}^r \gamma_i p(x_i) = \sum_{i=1}^r \frac{\gamma_i}{1 - \binom{0}{\mathbf{1}} x_i} x_i - \sum_{j=k+1}^d \left(\sum_{i=1}^r \frac{\gamma_i}{1 - \binom{0}{\mathbf{1}} x_i} (x_i)_j \right) e_j \quad (4.24)$$

where

$$\sum_{i=1}^r \frac{\gamma_i}{1 - \binom{0}{\mathbf{1}} x_i} - \sum_{j=k+1}^d \left(\sum_{i=1}^r \frac{\gamma_i}{1 - \binom{0}{\mathbf{1}} x_i} (x_i)_j \right) = \sum_{i=1}^r \frac{\gamma_i}{1 - \binom{0}{\mathbf{1}} x_i} \left(1 - \sum_{j=k+1}^d (x_i)_j \right) = \sum_{i=1}^r \gamma_i = 0$$

Since Q is a d -polytope, there is at least one column a_i of A that is not contained in the affine hull of the unit vectors, hence by the above calculation $p(a_i)$ cannot be contained in the affine hull of e_1, \dots, e_k , hence Q_k must be k -dimensional.

Now consider a facet $F = \text{conv}\{x_1, \dots, x_s, e_{k+1}, \dots, e_d\}$ of Q , given by some hyperplane with normal vector v . This means that for some $\mu \in \mathbb{R}$, $v^\top x \leq \mu$ for every

$x \in Q$, with equality exactly for $x \in F$. By scaling this inequality we can assume that $\mu = 1$ (if μ is negative, the direction of the inequality is reversed, but our argument still works). Since $e_i \in F$ for $i > k$, we get that $v_i = 1$ for those i .

We claim that F_k is defined as a face of Q_k by the hyperplane with normal vector $v(k) = (v_1, \dots, v_k, 0, \dots, 0)$. We need to prove that all vertices x of Q_k satisfy $v(k)^\top x \leq 1$, with equality if $x \in F_k$, and strict inequality otherwise. This is obvious for the unit vectors e_1, \dots, e_k ; for any column $a_i = (\alpha_1, \dots, \alpha_d)$ of A , the inequality $v^\top a_i < 1$ implies

$$v(k)^\top p(a) = \frac{\sum_{j=1}^k v_j \alpha_j}{1 - \sum_{j=k+1}^d \alpha_j} = \frac{\sum_{j=1}^k v_j \alpha_j}{1 - \sum_{j=k+1}^d v_j \alpha_j} \leq 1 \quad (4.25)$$

Since $\begin{pmatrix} 0 \\ \mathbf{1} \end{pmatrix}^\top a_i < 1$ for all columns a_i of A , the last inequality is strict if $v^\top a_i < 1$, and becomes an equality if $v^\top a_i = 1$. So F_k is indeed a face of Q_k , and since F had at least d extremal points, the argument at the beginning of the proof implies that F_k must have had at least k affinely independent points. Hence F_k is a proper face of maximal dimension, i.e. a facet. \square

The following Lemma provides a useful projective transformation with which we can influence the shape of the normals on a given polytope.

Lemma 4.24. *Consider a d -polytope Q such that $\mathbf{1}^\top x \geq 1$ for all $x \in Q$ and $F_0 = \{x \in Q \mid \mathbf{1}^\top x = 1\}$ is a facet of Q . Let F be another facet that can be written in the form $F = \{x \in Q \mid v^\top x = 1\}$ where $v^\top x \leq 1$ for all $x \in Q$. Let $0 \leq \mu < 1$, and consider the projective transformation*

$$\tau(x) = \frac{1}{\mu(\mathbf{1}^\top x) + 1 - \mu} \cdot x \quad (4.26)$$

Then in the transformed polytope $Q' = \tau(Q)$, the facet $\tau(F)$ has normal

$$w = \mu \mathbf{1} + (1 - \mu) \cdot v \quad (4.27)$$

Any point on the hyperplane $\{x \mid \mathbf{1}^\top x = 1\}$, which contains F_0 , is unchanged under τ .

Proof. The inverse projection map τ^{-1} that maps Q' to Q is given by

$$\tau^{-1}(y) = \frac{1 - \mu}{1 - \mu(\mathbf{1}^\top y)} \cdot y$$

For $x \in Q$ and hence $y = \tau(x) \in Q'$, we want that $v^\top x \leq 1$ is equivalent to $w^\top y \leq 1$. The latter inequality states $w^\top x / (\mu(\mathbf{1}^\top x) + 1 - \mu) \leq 1$, or $(w - \mu\mathbf{1})^\top x \leq 1 - \mu$. This follows from $w - \mu\mathbf{1} = (1 - \mu) \cdot v$, that is, when (4.27) holds, which also implies that $v^\top x = 1$ if and only if $w^\top y = 1$. \square

Proof of Proposition 4.15. We can now prove Proposition 4.15 for the case of non-disjoint facets, i.e. for $k < d$. As explained earlier, we would like to project the “disjoint” part of the polytope into a lower dimensional space, using the projection defined in (4.23). For this projection to be defined on the polytope, we need that the normal $\begin{pmatrix} 1 \\ \mathbf{0} \end{pmatrix} \alpha + \begin{pmatrix} \mathbf{0} \\ \mathbf{1} \end{pmatrix}$ on the facet F satisfies $\alpha > 0$. This can be achieved by applying the transformation τ in (4.26) for some μ close to 1. Since from now on we are only going to consider the transformed polytope, we denote the transformed polytope again by P^Δ , the transformed vertices by I, C, C' (where I is still the identity matrix), and the transformed completely labelled facets by F and F_0 . We will add suitable vertices to the transformed polytope to make the transformed facets unique. All we need to take care of is that at the end of our construction, we will be able to re-transform the extended polytope using τ^{-1} (this restriction is quite significant; it will force us to add a quadratically growing number of vertices, instead of the linearly growing number in the previous section).

Denote by p the projection defined in (4.23). After the transformation τ , the vertices c_i, c'_i of P^Δ still have positive entries, and the normal on the transformed facet F is of the form $\begin{pmatrix} 1 \\ \mathbf{0} \end{pmatrix} \alpha + \begin{pmatrix} \mathbf{0} \\ \mathbf{1} \end{pmatrix}$ for some $\alpha \in (0, 1)$. Hence for any column c of $[C \ C']$, we have that

$$c^\top \begin{pmatrix} \mathbf{0} \\ \mathbf{1} \end{pmatrix} < c^\top \left(\begin{pmatrix} 1 \\ \mathbf{0} \end{pmatrix} \alpha + \begin{pmatrix} \mathbf{0} \\ \mathbf{1} \end{pmatrix} \right) \leq 1$$

which implies that the projection p is defined for c . Consider the labelled k -polytope Q spanned by e_1, \dots, e_k and $p(c_1), \dots, p(c_k), p(c'_1), \dots, p(c'_n)$, which lives in the subspace of \mathbb{R}^d spanned by the first k unit vectors. The vertices e_i and $p(c_i)$ have label i (for $1 \leq i \leq k$); the vertices $p(c'_1), \dots, p(c'_n)$ might have labels in $\{k+1, \dots, d\}$, but these labels are irrelevant since those vertices will vanish when we add new labelled vertices later. By Lemma 4.23, the polytope Q has two disjoint completely labelled facets

$$G_0 = \text{conv}\{e_1, \dots, e_k\}, G = \text{conv}\{p(c_1), \dots, p(c_k)\}$$

Chapter 4. Index and Uniqueness of Symmetric Equilibria

Observe that since the normal on F is of the form $\begin{pmatrix} 1 \\ 0 \end{pmatrix} \alpha + \begin{pmatrix} 0 \\ 1 \end{pmatrix}$, we have

$$\alpha(c_i^\top \begin{pmatrix} 1 \\ 0 \end{pmatrix}) + c_i^\top \begin{pmatrix} 0 \\ 1 \end{pmatrix} = 1 \quad (4.28)$$

for $1 \leq i \leq k$. By definition of p , we get for $1 \leq i \leq k$

$$p(c_i)^\top \mathbf{1} = \frac{c_i^\top \begin{pmatrix} 1 \\ 0 \end{pmatrix}}{1 - \begin{pmatrix} 0 \\ 1 \end{pmatrix}^\top c_i} \stackrel{(4.28)}{=} 1/\alpha > 1 \quad (4.29)$$

hence the facets G and G_0 are parallel in k -dimensional space. Also, since

$$\det[p(c_1) \cdots p(c_k)] = t \cdot \det[c_1 \cdots c_k \ e_{k+1} \cdots e_d]$$

for some positive real number t , we get that the orientation of G is the same as the orientation of F , thus positive. Hence we can apply Proposition 4.18 to the polytope Q with completely labelled facets G and G_0 , and add vertices $y_1, \dots, y_{2k} \in \mathbb{R}^k \times \{0, \dots, 0\}$ with labels in $\{1, \dots, k\}$ to Q such that the only completely labelled facets of $Q_{\text{ext}} = \text{conv}(Q \cup \{y_1, \dots, y_{2k}\})$ are G and G_0 .

Assume that we could construct from the set $Y = \{y_1, \dots, y_{2k}\}$ of added vertices a new set of labelled points X that satisfy the following conditions:

- (i) the vertices C' of the polytope P^Δ are caught in the relative interior of the convex hull of the new points X and the last $d - k$ unit vectors, i.e.

$$C' \subset \text{relint}(\text{conv}[X, e_{k+1} \cdots e_d])$$

where by abuse of notation we write X for the matrix whose columns are given by the vectors in X .

- (ii) for each vector x in X , we get that $p(x)$ is in Y , and the labels of x and $p(x)$ agree.
- (iii) for each vector x in X , $\mathbf{1}^\top x < 1/\mu$ for the μ chosen at the beginning of the proof when we applied the projection τ from (4.26). This means that the inverse projection τ^{-1} is well-defined on each of the x .

We claim that those vertices X would do the trick for our original polytope. For this, we need to prove that $P_{\text{ext}}^\Delta = \text{conv}[I, C, C', X]$ has no completely labelled facets except for F and F_0 . Condition (i) above implies that for $1 \leq i \leq n$, c'_i vanishes in the relative interior of P_{ext}^Δ . By condition (ii), the labels of X are contained in $\{1, \dots, k\}$. In order for a facet \mathcal{F} to be completely labelled, it must then contain the vertices

e_{k+1}, \dots, e_d , and some other vertices $b_1, \dots, b_s \in X \cup \{c_1, \dots, c_k\} \cup \{e_1, \dots, e_k\}$ with labels $1, \dots, k$. By condition (ii) above and condition (d) of Proposition 4.18, each of these vertices is mapped by p to a vertex of Q_{ext} .

Then, Lemma 4.23 implies that $\mathcal{F}_k = \text{conv}\{p(b_1), \dots, p(b_s)\}$ is a completely labelled facet of the k -polytope Q_{ext} , i.e. \mathcal{F}_k has to be either G_0 or G . Since the only vertex of P_{ext}^Δ that is projected by p onto e_i is e_i , and the only vertex projected onto $p(c_i)$ is c_i , this implies that $b_i = e_i$ for all i , or that $b_i = c_i$ for all i , respectively. In the first case, we get $\mathcal{F} = F_0$, in the second, $\mathcal{F} = F$. Due to condition (iii) above, we can now re-transform P_{ext}^Δ using the transformation τ^{-1} given in (4.9). Since this re-transformation does not change the combinatorial structure of P_{ext}^Δ , we are done.

Hence all that we need to do is find points X satisfying conditions (i)-(iii) above. The original vertices Y will satisfy (i) and (ii), but in general not (iii). This is due to the following problem: By condition (c) of Proposition 4.18, we can choose a small positive ε such that every y_i in the “first” set $\{y_1, \dots, y_k\}$ of added vertices satisfies $\mathbf{1}^\top y_i = 1 + \varepsilon$, while every y_i in the “second” set $\{y_{k+1}, \dots, y_{2k}\}$ satisfies $\mathbf{1}^\top y_i = 1/\alpha - \varepsilon$. Since $\mu < 1$, we can assume that $\mathbf{1}^\top y_i < 1/\mu$ for the first set of vertices, but this inequality may not be true for the second set of added vertices. If we have $\mathbf{1}^\top y_i < 1/\mu$ for *all* added vertices we are done, by setting $X = Y$.

Otherwise, we revert to the following trick: Since the inverse transformation τ^{-1} is valid for the polytope P^Δ , we can conclude that $P^\Delta \subset \{x \mid \mathbf{1}^\top x < 1/\mu\}$. Choose $\varepsilon > 0$ such that both P^Δ and y_1, \dots, y_k are contained in the open half-space $\{x \mid \mathbf{1}^\top x < \frac{1}{\mu} - \varepsilon\}$. Denote by H the hyperplane $\{x \mid \mathbf{1}^\top x = \frac{1}{\mu} - \varepsilon\}$, and by H^- the corresponding closed halfspace containing P^Δ . By condition (b) of Proposition 4.18, $p(c'_i)$ is contained in the relative interior of $\text{conv}(Y)$ for $1 \leq i \leq n$. This, together with positivity of c'_i , implies that c'_i is in the relative interior of the d -polytope $\text{conv}[Y e_{k+1} \cdots e_d]$.

It is useful to adapt the first set of vertices $\{y_1, \dots, y_k\}$ slightly: For $1 \leq i \leq k$, we replace y_i by $x_i = e_d + \rho(y_i - e_d)$, where ρ is chosen such that $x_i \in H$ (i.e. $\rho = \frac{1/\mu - \varepsilon - 1}{\mathbf{1}^\top y_i - 1}$). The parameter ρ is independent of the choice of i due to condition (c) of Proposition 4.18. Since $\rho > 1$, the point y_i is a convex combination of x_i and e_d , and we get that C' is in the relative interior of the convex polytope $M = \text{conv}\{x_1, \dots, x_k, y_{k+1}, \dots, y_{2k}, e_{k+1}, \dots, e_d\}$. By our choice of the hyperplane H

it also follows that C' is contained in the relative interior of the polytope $M^- = M \cap H^-$. We would like to use the vertices of the latter polytope M^- to complete our desired set of points X .

What are the vertices of M^- , expressed in terms of the vertices of M ? M^- keeps those vertices of M that are contained in H^- , loses those vertices of M that are in the opposite open halfspace $H^+ \setminus H$, and gets a new vertex on H for each edge of M from a vertex of M in $H^- \setminus H$ to a vertex of M in $H^+ \setminus H$. Since we replaced the vertices y_1, \dots, y_k by vertices x_1, \dots, x_k contained in H , the only vertices of M in $H^- \setminus H$ are the unit vectors e_1, \dots, e_d . Hence M has vertices $e_{k+1}, \dots, e_d, x_1, \dots, x_k$, and additionally gets a vertex at each intersection of an edge of M with the hyperplane H that has y_i as one endpoint and e_j as the other endpoint, for some $k+1 \leq i \leq 2k$ and $1 \leq j \leq d$. For each $k+1 \leq i \leq 2k$, denote by X_i the set of vertices of M that arise from an edge between y_i and one of the e_j . What is the cardinality of X_i ? By condition (c) of Proposition 4.18, for any $k+1 \leq i \leq 2k$, there is no edge from y_i to any of the vertices e_1, \dots, e_k . Hence the worst that can happen is that we get an edge for each pair (y_i, e_j) , where $k+1 \leq i \leq 2k$, and $k+1 \leq j \leq d$. This means that for each y_i , the cardinality of X_i is at most $d-k$.

Viewing X_i as matrices, let $X = [x_1 \cdots x_k X_{k+1} \cdots X_{2k}]$, where x_i inherits the label of y_i , and the columns of X_j inherit the label of y_j . By construction, this set satisfies condition (i) above. Condition (iii) is true since by construction, $M^- \subset H^- \subset \{x \mid \mathbf{1}^\top x < 1/\mu\}$. As for condition (ii), observe that for any $x \in \mathbb{R}^k \times \mathbf{0} \subset \mathbb{R}^d$, any $\rho > 0$ and $j \in \{k+1, \dots, d\}$

$$\begin{pmatrix} \mathbf{0} \\ \mathbf{1} \end{pmatrix}^\top (e_j + \rho(x - e_j)) = 1 + \rho \left(\begin{pmatrix} \mathbf{0} \\ \mathbf{1} \end{pmatrix}^\top x - 1 \right) = 1 - \rho < 1$$

hence the projection p defined in (4.23) is defined on that point, and $p(e_j + \rho(x - e_j)) = x$. The last equation, which is true only since the last $d-k$ coordinates of x vanish, implies condition (ii). Hence we have found the desired set $X = [x_1 \cdots x_k X_{k+1} \cdots X_d]$ of labelled vertices that need to be added to the polytope P^Δ to make F and F_0 the only completely labelled facets. The cardinality of this set is at most $k + k(d-k)$. □

4.10 Open questions

In the previous section we have seen that for a symmetric equilibrium of positive index, with support size $k < d$ in a $d \times d$ symmetric game, we might need to add $k + k(d - k)$ strategies to make the equilibrium unique. As a first improvement on this bound we would hope to achieve a linear bound in k . However, an even stronger bound might be suggested: It is quite obvious that any pure Nash equilibrium of some bimatrix game can be made unique by adding a single strategy for the column player (see Lemma 4.17). This raises the question whether it should in general suffice to add k strategies, where k is the size of the support.

Another open question concerns our P-matrix construction. In Theorem 4.12 we proved that for P-matrices that are permutation-similar to an upper triangular matrix, the corresponding “canonical P-matrix-prism” has only two completely labelled facets. This statement is certainly true for any positive P-matrix as well, as long as the corresponding prism is simplicial. Otherwise Proposition 4.9 could be used to construct a principal minor of negative determinant. However, it is unclear in how far Theorem 4.12 holds for general P-matrices.

Bibliography

- Avis, D., G. Rosenberg, R. Savani, and B. von Stengel (2010), Enumeration of Nash equilibria for two-player games. *Economic Theory* 42, DOI 10.1007/s00199-009-0452-2, appeared online February 2009.
- Balthasar, A. (2010), Equilibrium tracing in strategic-form games. *Economic Theory* 42, DOI 10.1007/s00199-009-0442-4, appeared online February 2009.
- Bárány, I., S. Vempala, and A. Vetta (2007), Nash equilibria in random games. *Random Structures and Algorithms*, 31 (4), 391–405.
- Cohen, D. I. A. (1967), On the Sperner lemma. *J. Combinatorial Theory* 2, 585–587.
- Cottle, R. W., J.-S. Pang, and R. E. Stone (1992), *The Linear Complementarity Problem*. Academic Press, San Diego.
- Dantzig, G. B. (1963), *Linear Programming and Extensions*. Princeton University Press, Princeton, N.J.
- Demichelis, S., and F. Germano (2000), On the indices of zeros of Nash fields. *Journal of Economic Theory* 94, 192–217.
- Dold, A. (1980), *Lectures on Algebraic Topology*. Springer-Verlag, Berlin.
- Eaves, B. C. (1971), The linear complementarity problem. *Management Science* 17, 612–634.
- Eaves, B. C., and H. Scarf (1976), The solution of systems of piecewise linear equations. *Mathematics of Operations Research* 1, 1–27.

- Gale, D., H. W. Kuhn and A. W. Tucker (1950), On symmetric games, *Annals of Mathematics Studies* 24, 81–87.
- Gallier, J. (2001), *Geometric methods and applications. For computer science and engineering*, Springer-Verlag, New York.
- Garcia, C. B., and W. I. Zangwill (1981), *Pathways to Solutions, Fixed Points, and Equilibria*. Prentice-Hall, Englewood Cliffs.
- Govindan, S., A. von Schemde, and B. von Stengel (2003), Symmetry and p-Stability. *International Journal of Game Theory* 32, 359–269.
- Govindan, S., and R. Wilson (1997), Equivalence and invariance of the index and degree of Nash equilibria. *Games and Economic Behavior* 21, 56–61.
- Govindan, S., and R. Wilson (1997), Uniqueness of the index for Nash equilibria of two-player games. *Economic Theory* 10, 541–549.
- Govindan, S., and R. Wilson (2003), A global Newton method to compute Nash equilibria. *Journal of Economic Theory* 10, 65-86.
- Govindan, S., and R. Wilson (2003), Supplement to: A global Newton method to compute Nash equilibria. Accessed online at www.nyu.edu/jet/supplementary.html.
- Govindan, S., and R. Wilson (2005), Essential equilibria. *Proceedings of the National Academy of Sciences of the USA* 102, 15706–15711.
- Grünbaum, B. (2003), *Convex Polytopes*, 2nd ed. Springer-Verlag, New York.
- Harsanyi, J.C. (1975), The tracing procedure: A Bayesian approach to defining a solution for n -person noncooperative games. *International Journal of Game Theory*, 4, 61-94.
- Harsanyi, J.C., and R. Selten (1988), *A general theory of equilibrium selection in games*. MIT press, Cambridge.
- Hatcher, A. (2001), *Algebraic Topology*. Cambridge University Press.
- Hauk, E., and S. Hurkens (2002), On forward induction and evolutionary and strategic stability. *Journal of Economic Theory* 106, 66–90.

- Herings, P. J.-J., and R. Peeters (2009), Homotopy methods to compute equilibria in game theory. *Economic Theory* 42, DOI 10.1007/s00199-009-0441-5, appeared online February 2009.
- Hofbauer, J. (2003), Some thoughts on sustainable/learnable equilibria. Paper presented at the 15th Italian Meeting on Game Theory and Applications, Urbino, Italy, July 9–12, 2003. Accessed online at <http://www.econ.uniurb.it/imgta/PlenaryLecture/Hofbauer.pdf>.
- Hofbauer, J., and K. Sigmund (1998), *Evolutionary Games and Population Dynamics*. Cambridge University Press, Cambridge.
- Jansen, M. J. M. (1981), Maximal Nash subsets for bimatrix games. *Naval Research Logistics Quarterly* 28, 147–152.
- Johnson, C. R., D. D. Olesky, and P. van den Driessche (2003), Matrix classes that generate all matrices with positive determinant. *SIAM Journal on Matrix Analysis and Applications* 25, 285–294.
- Kohlberg, E., and J.-F. Mertens (1986), On the strategic stability of equilibria. *Econometrica* 54, 1003–1037.
- Lemke, C. E. (1965), Bimatrix equilibrium points and mathematical programming. *Management Science* 11, 681–689.
- Lemke, C. E., and S. J. Grotzinger (1976), On generalizing Shapley’s index theory to labelled pseudomanifolds. *Mathematical Programming* 10, 245–262.
- Lemke, C. E., and J. T. Howson, Jr. (1964), Equilibrium points of bimatrix games. *Journal of the Society for Industrial and Applied Mathematics* 12, 413–423.
- McLennan, A., and R. Tourky (2007), Imitation games and computation. Discussion Papers Series 359, School of Economics, University of Queensland, Australia.
- Mertens, J.-F. (1989), Stable equilibria – a reformulation, Part I. *Mathematics of Operations Research* 14, 575–625.
- Mertens, J.-F. (1991), Stable equilibria – a reformulation, Part II. *Mathematics of Operations Research* 16, 694–753.

- Morris, W. D., Jr. (1994): Lemke paths on simple polytopes, *Mathematics of Operations Research* 19, 780–789.
- Myerson, R. B. (1997), Sustainable equilibria in culturally familiar games. In: *Understanding Strategic Interaction: Essays in Honor of Reinhard Selten*, eds. W. Albers et al., Springer, Heidelberg, 111–121.
- Nash, J. F. (1951), Non-cooperative games. *Annals of Mathematics* 54, 286–295.
- O’Neill, B. (1953), Essential sets and fixed points. *American Journal of Mathematics* 75, 497–509.
- Papadimitriou, C. H. (1994), On the complexity of the parity argument and other inefficient proofs of existence. *Journal of Computer and System Sciences* 48, 498–532.
- Quint, T., and M. Shubik (1997), A theorem on the number of Nash equilibria in a bimatrix game. *International Journal of Game Theory* 26, 353–359.
- Ritzberger, K. (1994), The theory of Normal form games from the differentiable viewpoint, *Int. J. Game Theory* 23, 207–236.
- Ritzberger, K. (2002), *Foundations of Non-Cooperative Game Theory*. Oxford University Press, Oxford.
- Savani, R. (2006), Finding Nash equilibria of bimatrix games. *PhD thesis*.
- Savani, R., and B. von Stengel (2006), Hard-to-solve bimatrix games. *Econometrica* 74, 397–429.
- Shapley, L. S. (1974), A note on the Lemke–Howson algorithm. *Mathematical Programming Study 1: Pivoting and Extensions*, 175–189.
- Smale, S. (1976), A convergent process of price adjustment and global Newton methods. *Journal of Mathematical Economics* 3, 107–120.
- Sperner, E. (1928), Neuer Beweis für die Invarianz der Dimensionszahl und des Gebietes. *Abhandlungen aus dem Mathematischen Seminar der Hamburgischen Universität* 6, 265–272.

- Todd, M. J. (1976), Orientation in complementary pivot algorithms. *Mathematics of Operations Research* 1, 54–66.
- van Damme, E. (1987), *Stability and Perfection of Nash Equilibria*. Springer-Verlag, Berlin.
- van den Elzen, A. H., and A. J. J. Talman (1991), A procedure for finding Nash equilibria in bi-matrix games. *ZOR – Methods and Models of Operations Research* 35, 27–43.
- van den Elzen, A. H., and A. J. J. Talman (1999), An algorithmic approach toward the tracing procedure for bi-matrix games. *Games and Economic Behavior* 28, 130–145.
- von Schemde, A. (2005), *Index and Stability in Bimatrix Games*. Lecture Notes in Economics and Mathematical Systems, Vol. 1853, Springer-Verlag, Berlin.
- von Schemde, A., and B. von Stengel (2008), Strategic characterization of the index of an equilibrium. In: *Symposium on Algorithmic Game Theory (SAGT) 2008*, eds. B. Monien and U.-P. Schroeder, Lecture Notes in Computer Science, Vol. 4997, Springer-Verlag, Berlin, 242–254.
- von Stengel, B. (1996), Computing equilibria for two-person games. *Technical Report 253*, Dept. of Computer Science, ETH Zürich.
- von Stengel, B. (1999), New maximal numbers of equilibria in bimatrix games. *Discrete and Computational Geometry* 21, 557–568.
- von Stengel, B. (2002), Computing equilibria for two-person games. Chapter 45, *Handbook of Game Theory, Vol. 3*, eds. R. J. Aumann and S. Hart, North-Holland, Amsterdam, 1723–1759.
- von Stengel, B. (2007), Equilibrium computation for two-player games in strategic and extensive form. Chapter 3 of *Algorithmic Game Theory*, eds. N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, Cambridge Univ. Press, Cambridge, 53–78.
- von Stengel, B., A. H. van den Elzen, and A. J. J. Talman (2002), Computing normal form perfect equilibria for extensive two-person games. *Econometrica* 70, 693–715.

Bibliography

Vorob'ev, N. N. (1958), Equilibrium points in bimatrix games. *Theory of Probability and its Applications* 3, 297–309.

Ziegler, G. M. (1995), *Lectures on Polytopes*. Springer-Verlag, New York.