Homotopy methods to compute equilibria in game theory

Citation for published version (APA):

Herings, P. J. J., & Peeters, R. J. A. P. (2006). *Homotopy methods to compute equilibria in game theory*. METEOR, Maastricht University School of Business and Economics. METEOR Research Memorandum No. 046 <https://doi.org/10.26481/umamet.2006046>

Document status and date:

Published: 01/01/2006

DOI:

[10.26481/umamet.2006046](https://doi.org/10.26481/umamet.2006046)

Document Version:

Publisher's PDF, also known as Version of record

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

P. Jean-Jacques Herings, Ronald Peeters

Homotopy Methods to Compute Equilibria in
Game Theory

RM/06/046

JEL code: C62, C63, C72, C73

Maastricht research school of Economics
of TEchnology and ORganizations

Universiteit Maastricht
Faculty of Economics and Business Administration
P.O. Box 616
NL - 6200 MD Maastricht

phone : ++31 43 388 3830
fax : ++31 43 388 4873

Homotopy Methods to Compute Equilibria in Game Theory*

P. Jean-Jacques Herings

Ronald Peeters

October 2006

Abstract

This paper presents a complete survey of the use of homotopy methods in game theory. Homotopies allow for a robust computation of game-theoretic equilibria and their refinements. Homotopies are also suitable to compute equilibria that are selected by various selection theories. We present all relevant techniques underlying homotopy algorithms. We give detailed expositions of the Lemke-Howson algorithm and the Van den Elzen-Talman algorithm to compute Nash equilibria in 2-person games, and the Herings-Van den Elzen, Herings-Peeters, and McKelvey-Palfrey algorithms to compute Nash equilibria in general n -person games.

JEL Classification Codes: C62, C63, C72, C73.

Keywords: Homotopy, Equilibrium computation, Non-cooperative games, Nash Equilibrium.

1 Introduction

Many research fields have benefited greatly from the development of game theoretic tools in the twentieth century. The usual research strategy involves the modeling of the situation of interest as a game, followed by an analysis of the outcomes of the game that satisfy a particular solution concept, typically the one of Nash equilibrium or one of its refinements.

Since the analysis of Nash equilibria is key to the understanding of the situation at hand, it should not come as a surprise that much attention has been given to the development of methods by which Nash equilibria can be computed. We adhere to the view expressed in Judd (1997) that computational methods will serve an important role in the further development of economic theory, and in this case in the progress of game theory.

This survey provides an overview of a great variety of methods for Nash equilibrium computation that have been proposed. Although quite distinct in the specifics of the mathematical techniques used, all these methods share the property that they are guaranteed to converge to some Nash equilibrium. We argue that the homotopy idea is common to all these methods, and this idea should therefore be considered as the main driving force behind convergence.

*Department of Economics, Universiteit Maastricht, P.O. Box 616, 6200 MD, Maastricht, The Netherlands. E-mail: {P.Herings, R.Peeters}@algec.unimaas.nl. Both authors are financially supported by NWO.

The homotopy idea is also extremely helpful in obtaining a thorough understanding of the methods themselves.

Advantages of the use of homotopy methods include their numerical stability, their potential to be globally rather than only locally convergent, their ability to locate multiple solutions, and the insight they provide into properties of the solutions. Favorable by-products are that homotopies can be used for proving existence of an equilibrium and for the generic oddness of the number of equilibria.

The homotopy idea is very simple. It consists of, first, the problem of interest, second an artificial problem that can easily be solved, and third a continuous transformation of the easy to solve artificial problem into the problem of interest. The homotopy idea then consists of solving the easy artificial problem first and then using this solution to solve the transformations of the easy problem until finally the problem of interest has been solved. Section 2 explains why a deep mathematical result known as Browder's fixed point theorem (Browder, 1960), see also the extension by Mas-Colell (1974), is the fundamental reason that this approach works under very general assumptions.

We will refer to the easy to solve problem as the starting problem, the problem of interest as the terminal problem, and the problems created by the continuous transformation as the intermediate problems. The choice of different starting problems and the choice of different transformations of the starting problem to the terminal problem, creates different homotopy algorithms. We will explain all algorithms in these terms.

Our survey consists of two main parts. Part 1 (Sections 4 and 5) are devoted to the study of bi-matrix games, whereas Part 2 (Sections 6, 7, 8 and 9) treats general n -person games. The reason for this distinction is that the linearity present in bimatrix games makes it possible to find exact solutions for starting, intermediate and terminal problems, whereas general n -person games require some approximation method to find solutions.

Section 4 starts with a presentation of what is probably the best-known algorithm for the computation of Nash equilibria, the Lemke-Howson algorithm. This algorithm finds a Nash equilibrium by solving a related linear complementarity problem. A generalization of the Lemke-Howson algorithm is provided by the algorithm of van den Elzen and Talman (1991), described in Section 5. The latter algorithm allows for an arbitrarily chosen starting point.

Section 6 presents an overview of the use of homotopy methods in general n -person games. The algorithms of Sections 4 and 5 can be generalized to n -person games. For the Lemke-Howson procedure, such a generalization is given in Rosenmüller (1971) and Wilson (1971). Herings and van den Elzen (2002) extend the algorithm of van den Elzen and Talman (1991) to the n -person case. A difficulty to compute Nash equilibria in n -person games is posed by the non-linearity of the terminal problem. This makes it impossible to solve the intermediate problems exactly. Herings and van den Elzen (2002) tackle this problem by approximating the intermediate problems by piecewise linear ones, which can be solved exactly. Section 7 discusses the simplicial techniques that are needed for this approach. Section 8 presents the

algorithm of Herings and Peeters (2001). This algorithm exploits the differentiability present in game theoretic problems. The techniques needed for this approach are presented in detail.

The Herings-Van den Elzen and the Herings-Peeters algorithms, as well as the Van den Elzen-Talman algorithm, compute the Nash equilibrium selected by the equilibrium selection theory of Harsanyi and Selten (1988). Another algorithm with a nice game theoretic interpretation is the one related to the quantal-response equilibrium of McKelvey and Palfrey (1995). This approach is explained in detail in Section 9. Section 10 discusses how the methods presented in the other sections can be extended to extensive form games and dynamics games. Section 11 concludes.

2 Homotopy Methods

In topology, two continuous functions from one topological space to another are called homotopic if one can be “continuously deformed” into the other. Such a deformation is called a homotopy between the two functions. Formally, a homotopy between two continuous functions f and g from a topological space X to a topological space Y is defined to be a continuous function $H : [0, 1] \times X \rightarrow Y$ such that, for all points x in X , $H(0, x) = f(x)$ and $H(1, x) = g(x)$. If we think of the homotopy-parameter in $[0, 1]$ as “time”, then H describes a “continuous deformation” of f into g : at time 0 we have the function f , at time 1 we have the function g .

A properly defined homotopy can be used to solve for the zeros of the function g starting from the zeros of the function f and gradually solve the nearby systems for increasing values of the homotopy parameter. These types of procedures are called homotopy methods. The following mathematical result due to Browder (1960) delivers the key insight for why homotopy methods converge under very general circumstances.

THEOREM 2.1 *Let S be a non-empty, compact, convex subset of $\mathbb{R}^m$ and let $H : [0, 1] \times S \rightarrow S$ be a continuous function. Then the set of fixed points, $F_H = \{(\lambda, s) \in [0, 1] \times S \mid s = H(\lambda, s)\}$ contains a connected set, F_H^c , such that $(\{0\} \times S) \cap F_H^c \neq \emptyset$ and $(\{1\} \times S) \cap F_H^c \neq \emptyset$.*

The set of fixed points F_H of H contains a connected subset F_H^c that intersects the two extremes: $\{0\} \times S$ and $\{1\} \times S$. This result suggests the following algorithm. Represent the problem for which a solution has to be computed, the terminal problem, as a fixed point problem, i.e. formulate a function $H : \{1\} \times S \rightarrow S$ such that a fixed point of H yields a solution to the terminal problem. Formulate the starting problem as a fixed point problem, an artificial function $H : \{0\} \times S \rightarrow S$ with usually a unique fixed point, say s^0 , that can be computed easily. Define $H : (0, 1) \times S \rightarrow S$ in any way that makes H continuous on $[0, 1] \times S$. By Theorem 2.1, the point $(0, s^0)$ is connected by F_H^c to a fixed point $(1, s^1)$ of H . Homotopy methods are methods to follow the connected set F_H^c . A number of techniques is available to do so, depending on the properties of F_H^c .

The following result, due to Mas-Colell (1974), shows that Theorem 2.1 can be generalized to the case of upper hemi-continuous correspondences.

THEOREM 2.2 *Let S be a non-empty, compact, convex subset of $\mathbb{R}^m$ and let $H : [0, 1] \times S \rightarrow S$ be an upper hemi-continuous correspondence that is non-empty, compact, and convex-valued. Then the set of fixed points, $F_H = \{(\lambda, s) \in [0, 1] \times S \mid s \in H(\lambda, s)\}$ contains a connected set, F_H^c , such that $(\{0\} \times S) \cap F_H^c \neq \emptyset$ and $(\{1\} \times S) \cap F_H^c \neq \emptyset$.*

Since the fixed point problem related to H is a problem with $m + 1$ free variables and m equations, it is usually possible to formulate appropriate regularity conditions for which the solution is a compact, piecewise differentiable 1-dimensional manifold, i.e. a finite collections of arcs and loops. If the starting problem is constructed to have a unique fixed point, then there is a unique arc (also referred to as path or homotopy path) from this unique fixed point to a fixed point in $\{1\} \times S$. We will also present methods that work without such regularity conditions.

There are two fundamental methods of numerically tracing the homotopy path: predictor-corrector methods, and simplicial methods. Predictor-corrector methods approximately follow exact solution curves, whereas simplicial methods exactly follow approximate solution curves. For some problems the homotopy path is piecewise linear. Numerically tracking is not needed, and it is possible to compute the homotopy path exactly by a finite sequence of linear programming steps. This applies in particular to bimatrix games. After introducing general notation for games in Section 3, we will turn to bimatrix games in Sections 4 and 5. For a detailed overview on path-tracking methods the reader is referred to Garcia and Zangwill (1981) and Allgower and Georg (1990). For an overview of homotopy methods in economic theory, see Eaves and Schmedders (1996).

3 Notation

A finite n -person noncooperative game in normal form is a tuple $\Gamma = \langle N, \{S^i\}_{i \in N}, \{u^i\}_{i \in N} \rangle$, where $N = \{1, \dots, n\}$ is the finite set of players, $S^i = \{s_1^i, \dots, s_{m_i}^i\}$ is the finite set of actions that player i has at his disposal, and where u^i is a real valued function on the set of all possible action combinations $S = \times_{i \in N} S^i$. A mixed action of player i is a probability distribution on S^i . Thereto we define Σ^i as the set of all probability distributions over S^i . For $\sigma^i \in \Sigma^i$, the probability assigned to pure action s_j^i is given by σ_j^i . The payoff function is extended multi-linearly to the set of all mixed action combinations $\Sigma = \times_{i \in N} \Sigma^i$. Given a mixed action combination $\sigma \in \Sigma$ and a mixed action $\bar{\sigma}^i \in \Sigma^i$, we denote by $(\sigma^{-i}, \bar{\sigma}^i)$ the mixed action combination that results from replacing σ^i by $\bar{\sigma}^i$. A mixed strategy combination $\sigma \in \Sigma$ is said to be a Nash equilibrium of game Γ if σ^i is a best response against σ^{-i} for all $i \in N$. The set of Nash equilibria of game Γ is denoted by $\text{NE}(\Gamma)$.

We define the disjoint union of the pure strategies over all players by $S^* = \bigcup_{i \in N} S^i$ with $m^* = \sum_{i \in N} m^i$ as the cardinality of this set. A subset B^* of S^* is called admissible if $B^* \cap S^i \neq \emptyset$ for all $i \in N$. An admissible subset B^* of S^* contains at least one pure strategy for each player.

The case $n = 2$ is referred to as a bimatrix game. A bimatrix game can be represented by a pair of $m^1 \times m^2$ matrices (A^1, A^2) , where $u^1(\sigma^1, \sigma^2) = (\sigma^1)^\top A^1 \sigma^2$ and $u^2(\sigma^1, \sigma^2) = (\sigma^1)^\top A^2 \sigma^2$.

The vector e^i is the column-vector of length m^i for which each element is equal to 1, e_j^i is the column-vector of length m^i for which the j th element is equal to 1 and all other elements 0, and 0^i is the column-vector of length m^i for which each element is equal to 0.

For a matrix M , we define by M_i the i th row and by $M_{\cdot j}$ the j th column of the matrix M .

4 Bimatrix Games: The Lemke-Howson Algorithm

The problem of finding a Nash equilibrium in bimatrix games is equivalent to solving a linear complementarity problem. As a result mathematical programming techniques that are developed for solving linear complementarity problems are applicable to solve bimatrix games. For a recent overview on the computation of equilibria for bimatrix games, we refer the interested reader to von Stengel (2002).

The first path following algorithm to solve a nondegenerate bimatrix game has been developed by Lemke and Howson (1964). Their method finds a Nash equilibrium by solving a related linear complementarity problem.

THEOREM 4.1 (LEMKE AND HOWSON, 1964) *A strategy pair $(\sigma^1, \sigma^2) \in \Sigma$ constitutes a Nash equilibrium of the bimatrix game (A^1, A^2) if and only if there exists a pair $(\mu^1, \mu^2) \in \mathbb{R} \times \mathbb{R}$ such that*

$$\begin{aligned} (\sigma^1)^\top [(e^1)^\top \mu^1 - A^1 \sigma^2] &= 0 & (\sigma^2)^\top [(e^2)^\top \mu^2 - (A^2)^\top \sigma^1] &= 0 \\ \sigma^1 &\geq 0^1 & \sigma^2 &\geq 0^2 \\ (e^1)^\top \mu^1 - A^1 \sigma^2 &\geq 0^1 & (e^2)^\top \mu^2 - (A^2)^\top \sigma^1 &\geq 0^2 \\ (e^1)^\top \sigma^1 - 1 &= 0 & (e^2)^\top \sigma^2 - 1 &= 0. \end{aligned} \tag{1}$$

The pair (μ^1, μ^2) is the pair of equilibrium payoffs corresponding to the Nash equilibrium (σ^1, σ^2) . The second and fourth lines determine that σ^1 and σ^2 are mixed actions, with probabilities assigned to pure actions being non-negative and summing up to one. The third line makes sure that there is no pure action of player 1 (2) that gives a higher payoff than μ^1 (μ^2) when player 2 (1) plays σ^2 (σ^1). According to the first line each pure action s_j^1 (s_j^2) is played with zero probability or leads to the payoff μ^1 (μ^2) and is therefore optimal. It allows for an optimal action to be played with probability zero.

Theorem 4.1 states that the problem of finding a Nash equilibrium of the bimatrix game (A^1, A^2) is equivalent to finding a solution to the linear complementarity problem (LCP) of Eq. (1). Like von Stengel (2002), we follow Shapley's (1974) exposition of the Lemke-Howson algorithm. The algorithm is mainly based on the property that in an equilibrium each pure action is played with probability zero or is a best response (or both).

Any mixed action σ^i is assigned a set of labels in S^* . In particular, σ^i is labeled by the pure actions that are played with probability zero and the pure actions of the opponent that are best responses to it:

$$\begin{aligned} L^1(\sigma^1) &= \{s_j^1 \in S^1 \mid \sigma_j^1 = 0\} \cup \{s_j^2 \in S^2 \mid (A_{\cdot j}^2)^\top \sigma^1 \geq (A_{\cdot \ell}^2)^\top \sigma^1 \text{ } (s_\ell^2 \in S^2)\}, \\ L^2(\sigma^2) &= \{s_j^2 \in S^2 \mid \sigma_j^2 = 0\} \cup \{s_j^1 \in S^1 \mid A_{\cdot j}^1 \sigma^2 \geq A_{\cdot \ell}^1 \sigma^2 \text{ } (s_\ell^1 \in S^1)\}. \end{aligned} \quad (2)$$

A profile of mixed actions $\sigma = (\sigma^1, \sigma^2)$ is labeled by the union of the labels of σ^1 and σ^2 :

$$L(\sigma) = L^1(\sigma^1) \cup L^2(\sigma^2). \quad (3)$$

A profile of mixed actions σ is called *completely labeled* if $L(\sigma) = S^*$ and called *s_j^i -almost completely labeled* if $L(\sigma) = S^* \setminus \{s_j^i\}$. The following result can easily be shown.

THEOREM 4.2 *A mixed action pair $(\sigma^1, \sigma^2) \in \Sigma$ constitutes a Nash equilibrium of the bimatrix game (A^1, A^2) if and only if it is completely labeled.*

A bimatrix game with the property that any $\sigma^i \in \Sigma^i$ has at most m^i labels is called *nondegenerate*. In this section, we will restrict the analysis to nondegenerate games. Wilson (1992) amended the Lemke-Howson algorithm to handle non-generic games and, moreover, to select a simply stable equilibrium.

THEOREM 4.3 *In a nondegenerate bimatrix game (A^1, A^2) only finitely many points in Σ^1 have m^1 labels and only finitely many points in Σ^2 have m^2 labels.*

It follows as a result of this theorem that in a nondegenerate bimatrix game there are finitely many points in $\Sigma^1 \times \Sigma^2$ that have $m^1 + m^2$ labels. Consequently, there are finitely many points that are completely labeled and hence finitely many Nash equilibria.

Consider a nondegenerate bimatrix game (A^1, A^2) and let G^i (for $i = 1, 2$) be the undirected graph whose vertices are the points from Σ^i with m^i labels, and an additional vertex 0^i having all actions of player i as label ($L^i(0^i) = S^i$). Any two vertices of G^i are joined by an edge if they have exactly $m^i - 1$ labels in common (and thus differ in one label precisely).

Let $G = G^1 \times G^2$ be the product graph with the vertices being all pairs of vertices of G^1 and G^2 and the edges being all vertex-edge pairs with vertices coming from G^1 and edges from G^2 or vice versa. In line with the definitions above, an edge of G is called *s_j^i -almost completely labeled* if the two vertices that are connected by it are *s_j^i -almost completely labeled*.

A first observation is that for every s_j^i any completely labeled vertex (σ^1, σ^2) in G , that is all Nash equilibrium and the additional vertex $(0^1, 0^2)$, is adjacent to exactly one s_j^i -almost completely labeled vertex $(\tilde{\sigma}^1, \tilde{\sigma}^2)$. If s_j^i is a label of σ^1 , then σ^1 is adjacent to $\tilde{\sigma}^1$ in G^1 (sharing the same remaining $m^1 - 1$ labels) and $\sigma^2 = \tilde{\sigma}^2$. If s_j^i is a label of σ^2 , then σ^2 is adjacent to $\tilde{\sigma}^2$ in G^2 (sharing the same remaining $m^2 - 1$ labels) and $\sigma^1 = \tilde{\sigma}^1$. A second observation is that any s_j^i -almost completely labeled vertex (σ^1, σ^2) in G is adjacent to exactly two s_j^i -almost completely labeled vertices in G . A vertex (σ^1, σ^2) in G can only be s_j^i -almost completely labeled if σ^1 and σ^2 both have one other label s_ℓ^k in common. One neighbor of (σ^1, σ^2) is $(\tilde{\sigma}^1, \sigma^2)$ with $\tilde{\sigma}^1$ being a adjacent to σ^1 by a s_ℓ^k -almost completely labeled edge in G^1 . The other is $(\sigma^1, \tilde{\sigma}^2)$ with $\tilde{\sigma}^2$ being a adjacent to σ^2 by a s_ℓ^k -almost completely labeled edge in G^2 .

THEOREM 4.4 (LEMKE AND HOWSON, 1964; SHAPLEY, 1974) *Let (A^1, A^2) be a nondegenerate bimatrix game and s_j^i be a label in S^* . Then the set of s_j^i -almost completely labeled vertices and edges in G consists of disjoint paths and cycles. The end-points of the paths are the completely labeled vertices (the equilibria of the game) and the completely labeled vertex $(0^1, 0^2)$ (the artificial equilibrium). The number of Nash equilibria of the game is odd.*

The algorithm of Lemke-Howson starts in the artificial equilibrium $(0^1, 0^2)$ that has all labels. For given label s_j^i a path consisting of s_j^i -almost completely labeled edges and vertices is followed that terminates at a Nash equilibrium of the game. From a computational point of view, moving from one completely labeled edge to another, is equivalent to making a linear programming pivot step. We will explain this in more detail for the Van den Elzen-Talman algorithm presented in Section 5.

The essential reason that these pivoting methods converge is at best explained as follows. Suppose there is a haunted house with a finite number of rooms. Each room has either zero, one or two doors. There is only one entrance to the haunted house through which you just came in. Since you are being haunted by a ghost there is no way back. The only way to escape from the ghost is to hide in a room with one door (to which ghosts have no entry). If the room you entered by entering the haunted house has one door, you are safe. If not, then there is precisely one door leading to another room to which you have to escape. The room you entered has at least one door (the one you entered). Again, either you are safe or there is another door through which you should escape. This procedure either ends in a room with one door only or continues ad infinitum. The latter can however not happen. Namely, by finiteness of the number of rooms there is a moment that a room once visited is entered for the second time. Since it is impossible to enter this room by the two doors already used (this would imply that the room you are coming from was already visited twice), this room must have a third door. This is not possible, and hence it is impossible to wander ad infinitum. You can thus be sure to end up in a room having one door only. Using this door-in door-out

principle we find convergence to a room with one door. In general this principle can be used to show existence and oddness of the number of rooms with one door.

Consider the game of Figure 1. This game possesses a unique Nash equilibrium, the mixed action profile $(\bar{\sigma}^1, \bar{\sigma}^2) = ((\frac{2}{3}, \frac{1}{3}), (\frac{3}{4}, \frac{1}{4}))$.

	s_1^2	s_2^2
s_1^1	2,2	1,4
s_2^1	1,4	4,0

Figure 1: Example.

Figure 2 contains the graphs G^1 and G^2 for this game. The vertices are indicated by the (mixed) actions, the labels by the pure actions with quotation marks. The path of s_1^1 -almost completely labeled vertices and edges connects the artificial equilibrium $(0^1, 0^2)$ via the vertices $(s_1^1, 0^2)$, (s_1^1, s_2^2) , $(\bar{\sigma}^1, s_2^2)$ to the equilibrium $(\bar{\sigma}^1, \bar{\sigma}^2)$.

Figure 2: Example Lemke-Howson.

Though perhaps not straightforward from the exposition above, the Lemke-Howson method is a homotopy method. This can be seen as follows. In general, along the path of an s_j^i -almost completely labeled path starting from the artificial equilibrium, the first vertex attained is the profile composed by pure action s_j^i and the opponent's best response to it. This profile would constitute an equilibrium if s_j^i would be a best response in return. Generally this is not the case, but by giving player i a sufficiently large bonus for playing the pure action s_j^i , the first action profile could be an equilibrium. Let α be a bonus sufficiently large to make any pure action a dominant action once this bonus is given for the use of this action, for instance $\alpha = \max_{i,\sigma} u^i(\sigma) - \min_{i,\sigma} u^i(\sigma) + \varepsilon$. According to Wilson (1992) the path of s_j^i -almost completely labeled vertices and edges corresponds to a homotopy path of a homotopy involving a gradually diminishing bonus for the use of action s_j^i starting from a bonus of size α . To be more precise, the gradually diminishing bonus defines a series of bimatrix games $\Gamma(t)$ with $t \in [0, 1]$ and payoff function $v^i(t; s)$ such that $v^k(t; s) = u^k(s) + (1 - t)\alpha$ if $s^k = s_j^i$ and $v^k(t; s) = u^k(s)$ otherwise. Now, by construction, the first action profile on the path is an equilibrium of the game $\Gamma(0)$. By gradually increasing t from 0 to 1, the bonus on s_j^i is

gradually decreased from α to 0 where the original game is reached ($\Gamma = \Gamma(1)$). During this procedure a path of Nash equilibria is generated that starts in the first action profile in the Lemke-Howson procedure and ends in the Nash equilibrium selected by the Lemke-Howson procedure.

To obtain an explicit formulation as a homotopy, it suffices to define the correspondence $H : [0, 1] \times \Sigma \rightarrow \Sigma$ by

$$H(t, \sigma) = \prod_{i \in N} \beta^i(t; \sigma),$$

where $\beta^i(t; \sigma) = \operatorname{argmax}_{\bar{\sigma}^i \in \Sigma^i} v^i(t; \sigma^{-i}, \bar{\sigma}^i)$ is the best response correspondence of player i in game $\Gamma(t)$. The set of strategies of the resulting homotopy path coincides with the set of strategies generated by the Lemke-Howson algorithm. Convergence of the Lemke-Howson algorithm for a nondegenerate game now follows easily from Theorem 2.2.

The homotopy H of the previous paragraph is a correspondence. Using the techniques of Herings (2000) and Geanakoplos (2003), it is possible to define a continuous homotopy function with exactly the same homotopy path.

An alternative way to obtain a homotopy whose homotopy path is equivalent to the Lemke-Howson path, is by keeping the utility function of the terminal problem, and manipulating the strategy space. In particular, for given s_j^i , we define $\Sigma^i(t) = (1 - t)\{s_j^i\} + t\Sigma^i$ and $\Sigma^{-i}(t) = \Sigma^{-i}$. We define the correspondence $H : [0, 1] \times \Sigma \rightarrow \Sigma$ by

$$H(t, \sigma) = \prod_{i \in N} \beta_{\Sigma^i(t)}^i(\sigma),$$

where $\beta_{\Sigma^i(t)}^i(\sigma) = \operatorname{argmax}_{\bar{\sigma}^i \in \Sigma^i(t)} u^i(\sigma^{-i}, \bar{\sigma}^i)$ is the best response correspondence of player i in game Γ with restricted strategy set $\Sigma^i(t)$. The set of strategies of the resulting homotopy path coincides with the set of strategies generated by the Lemke-Howson algorithm.

5 Bimatrix Games: The Van den Elzen-Talman Algorithm

The Lemke-Howson algorithm always starts in a pure strategy profile, where for one player a pure strategy is randomly chosen, and the other player uses a best response. Van den Elzen and Talman (1991) developed an algorithm similar to the Lemke-Howson algorithm that allows for an arbitrary mixed strategy profile as the starting point. As before, nondegenerate bimatrix games are considered.

The advantage of an arbitrary starting point is twofold. First, often some reasonable guess of the actions that should be played in an equilibrium is available. Then it is natural to take as a starting point of the algorithm a strategy that puts only weight on such actions. Secondly, if there is an interest in detecting whether a given game has multiple Nash equilibria, or there is a desire to compute several Nash equilibria if there exist multiple ones, the flexibility of the starting point is a desirable feature.

Let a starting vector $\nu \in \Sigma$ be given. For an admissible subset B^* of S^* , Van den Elzen and Talman (1991) consider the following system of equations.

$$\begin{aligned}
\sigma_j^i &= (1-t)\nu_j^i + \tau_j^i & (s_j^i \in S^*) \\
u^i(\sigma^{-i}, s_j^i) + \lambda_j^i &= \mu^i & (s_j^i \in S^*) \\
\sum_{s_j^i \in S^i} \tau_j^i &= t & (i \in N) \\
\tau_j^i &= 0 & (s_j^i \notin B^*) \\
\lambda_j^i &= 0 & (s_j^i \in B^*)
\end{aligned} \tag{4}$$

with $t \in [0, 1]$, $\tau^i, \lambda^i \in \mathbb{R}_+^m$ and $\mu^i \in \mathbb{R}$. Suppose $(t, \sigma, \tau, \lambda, \mu)$ is a solution to (4) (notice that not all B^* admit solutions in general). The variable μ^i represents the maximum payoff that player i can get against strategy σ^{-i} . The pure actions in B^* are all best responses. The λ_j^i -s corresponding to best responses are therefore equal to zero. For actions not in B^* , the corresponding variable τ_j^i is equal to zero.

For bimatrix games the system above is linear and the second line of it can be replaced by

$$\begin{aligned}
A_j^1 \cdot \sigma^2 + \lambda_j^1 &= \mu^1 & (s_j^1 \in S^1) \\
(\sigma^1)^\top A_{\cdot j}^2 + \lambda_j^2 &= \mu^2 & (s_j^2 \in S^2).
\end{aligned} \tag{5}$$

By substitution of the first line of Equation (4) into these new equations we end up with the following system of $2(m^1 + m^2 + 1)$ linear equations and $2(m^1 + m^2 + 1) + 1$ variables.¹

$$\begin{aligned}
A_j^1 \cdot ((1-t)\nu^2 + \sum_{s_\ell^2 \in B^*} \tau_\ell^2 e_\ell^2) + \lambda_j^1 &= \mu^1 & (s_j^1 \in S^1) \\
((1-t)\nu^1 + \sum_{s_\ell^1 \in B^*} \tau_\ell^1 e_\ell^1)^\top A_{\cdot j}^2 + \lambda_j^2 &= \mu^2 & (s_j^2 \in S^2) \\
\sum_{s_j^i \in S^i} \tau_j^i &= t & (i \in N) \\
\tau_j^i &= 0 & (s_j^i \notin B^*) \\
\lambda_j^i &= 0 & (s_j^i \in B^*).
\end{aligned} \tag{6}$$

Or, in matrix-vector notation,

$$\begin{aligned}
(1-t) \begin{pmatrix} A^1 \nu^2 \\ (A^2)^\top \nu^1 \\ 1 \\ 1 \end{pmatrix} + \sum_{s_\ell^1 \in B^*} \tau_\ell^1 \begin{pmatrix} 0^1 \\ (A_{\ell \cdot}^2)^\top \\ 1 \\ 0 \end{pmatrix} + \sum_{s_\ell^2 \in B^*} \tau_\ell^2 \begin{pmatrix} A_{\cdot \ell}^1 \\ 0^2 \\ 0 \\ 1 \end{pmatrix} \\
+ \sum_{s_\ell^1 \notin B^*} \lambda_\ell^1 \begin{pmatrix} e_\ell^1 \\ 0^2 \\ 0 \\ 0 \end{pmatrix} + \sum_{s_\ell^2 \notin B^*} \lambda_\ell^2 \begin{pmatrix} 0^1 \\ e_\ell^2 \\ 0 \\ 0 \end{pmatrix} - \mu^1 \begin{pmatrix} e^1 \\ 0^2 \\ 0 \\ 0 \end{pmatrix} - \mu^2 \begin{pmatrix} 0^1 \\ e^2 \\ 0 \\ 0 \end{pmatrix} &= \begin{pmatrix} 0^1 \\ 0^2 \\ 1 \\ 1 \end{pmatrix}.
\end{aligned} \tag{7}$$

Given the number of equations and unknowns, for each admissible set B^* a one-dimensional solution set is what one expects. Van den Elzen and Talman (1991) define a bimatrix game to be nondegenerate if the solution set is 1-dimensional for each admissible set B^* . As usual,

¹Notice that $\tau^i = \sum_{s_\ell^i \in B^*} \tau_\ell^i e_\ell^i$ and $\lambda^i = \sum_{s_\ell^i \notin B^*} \lambda_\ell^i e_\ell^i$.

the empty set is considered to have any dimension. This means that for each B^* the set of solutions to the corresponding system is an empty set or a compact line segment.

Let B^* be the set containing for each player i his best response against ν^{-i} . For nondegenerate games, this best response, and thereby B^* , is uniquely defined. It is easily verified that $t = 0$, $\tau = 0$, $\mu^i = u^i(\nu^{-i}, s_j^i)$ for $s_j^i \in B^*$, and $\lambda_j^i = \mu^i - u^i(\nu^{-i}, s_j^i)$, $s_j^i \in S^*$, is a solution to (7). Notice that λ_j^i is strictly positive for $s_j^i \notin B^*$ and τ_j^i is strictly positive for $s_j^i \in B^*$.

The idea is to keep B^* fixed and to increase t until one of the constraints would get violated, i.e. one of the λ -s or τ -s would become negative. Since the set of solutions is a line segment, increasing t leads from the solution just identified, to the other boundary point of the line segment. From a computational point of view, the operation just described corresponds to a linear programming pivot step.

If the other boundary point of the line segment corresponds to $t = 1$, the algorithm stops. It is easily verified that a solution to (7) with $t = 1$ yields a Nash equilibrium of the bimatrix game Γ . Otherwise, we have to adapt the set B^* such that the procedure can be continued. For the first line segment generated, the case where τ_j^i becomes negative for some $s_j^i \in B^*$ cannot occur. The end-point of the line segment is characterized by $\lambda_j^i = 0$ for some $s_j^i \notin B^*$. In this case s_j^i is added to the set B^* and the system of equations (7) is studied for $B^* \cup \{s_j^i\}$. The end-point of the previous line segment is the starting point of the line segment of solutions to (7) for $B^* \cup \{s_j^i\}$. The end-point of this line segment satisfies either $t = 1$, in which case a Nash equilibrium has been found, or $\tau_j^i = 0$ for some $s_{j'}^{i'} \in B^* \cup \{s_j^i\}$, in which case a new admissible set is defined as $B^* \cup \{s_j^i\} \setminus \{s_{j'}^{i'}\}$, or $\lambda_{j'}^{i'} = 0$ for some $s_{j'}^{i'} \notin B^* \cup \{s_j^i\}$, and a new admissible set is defined as $B^* \cup \{s_j^i\} \cup \{s_{j'}^{i'}\}$.

The procedure stops when t is equal to 1. The door-in door-out principle of Lemke and Howson guarantees that such will be the case after generating a finite number of admissible sets B^* . Hence starting from a solution corresponding to $(t, \sigma) = (0, \nu)$ by linear programming steps a piecewise linear path towards a Nash equilibrium is followed. The starting profile plays an important role during the procedure since all adjustments are made relative to this initial profile.

THEOREM 5.1 *Let (A^1, A^2) be a nondegenerate bimatrix game and let ν belong to the interior of Σ . Then the algorithm finds a perfect Nash equilibrium.*

Figure 3 displays the projection on Σ of the path for the game of Figure 1 for initial starting profile $\nu = ((\frac{4}{5}, \frac{1}{5}), (\frac{2}{5}, \frac{3}{5}))$. For player 1 and 2, the pure actions s_2^1 and s_2^2 are the best responses against the initial action profile ν . Therefore we start tracking the solutions to the system in (7) for $B^* = \{s_2^1, s_2^2\}$ starting at $t = 0$. We can increase t until t reaches $\frac{1}{6}$, the strategy profile σ equals $((\frac{2}{3}, \frac{1}{3}), (\frac{1}{3}, \frac{2}{3}))$, and λ_1^2 becomes 0. This means that at $t = \frac{1}{6}$ the pure action s_1^2 has become a best response for player 2 and we should continue tracking the system determined by B^* extended with s_1^2 . In the new system for $B^* = \{s_2^1, s_1^2, s_2^2\}$ we

can keep t at $\frac{1}{6}$ and increase τ_1^2 until τ_2^2 becomes 0 and action s_2^2 should be eliminated from the admissible subset B^* . The strategy profile reached equals $\sigma = ((\frac{2}{3}, \frac{1}{3}), (\frac{1}{2}, \frac{1}{2}))$. The next system to solve is therefore determined by $B^* = \{s_2^1, s_1^2\}$ and we move in the direction of (s_2^1, s_1^2) by increasing t , τ_2^1 , and τ_1^2 . At $t = \frac{7}{12}$, $\sigma = ((\frac{1}{3}, \frac{2}{3}), (\frac{3}{4}, \frac{1}{4}))$, λ_1^1 becomes 0 and s_1^1 should be included in B^* : $B^* = \{s_1^1, s_2^1, s_1^2\}$. Next the path stays at $t = \frac{7}{12}$ and τ_1^1 is increased. We can increase τ_1^1 until $\tau_1^1 = \frac{1}{3}$ (and $\tau_2^1 = \frac{1}{4}$) when λ_1^1 gets 0. The strategy profile reached equals $\sigma = ((\frac{1}{3}, \frac{2}{3}), (\frac{3}{4}, \frac{1}{4}))$, the Nash equilibrium of the game. We complete the algorithm by following the path for $B^* = \{s_1^1, s_2^1, s_1^2, s_2^2\}$ keeping σ at the equilibrium value, towards $t = 1$.

Figure 3: Example van den Elzen-Talman.

The procedure tracks a piecewise linear path starting at $(0, \nu) = (0, ((\frac{4}{5}, \frac{1}{5}), (\frac{2}{5}, \frac{3}{5})))$, proceeding to $(\frac{1}{6}, ((\frac{2}{3}, \frac{1}{3}), (\frac{1}{3}, \frac{2}{3})))$, $(\frac{1}{6}, ((\frac{2}{3}, \frac{1}{3}), (\frac{1}{2}, \frac{1}{2})))$, $(\frac{7}{12}, ((\frac{1}{3}, \frac{2}{3}), (\frac{3}{4}, \frac{1}{4})))$, and $(\frac{7}{12}, ((\frac{2}{3}, \frac{1}{3}), (\frac{3}{4}, \frac{1}{4})))$, and reaching $t = 1$ with the Nash equilibrium at $(1, \bar{\sigma}) = (1, ((\frac{2}{3}, \frac{1}{3}), (\frac{3}{4}, \frac{1}{4})))$. Clearly, we could have terminated the algorithm at $t = \frac{7}{12}$ when we found the Nash equilibrium. More generally, as soon as a set $B^* = S^*$ is reached, a fully mixed Nash equilibrium has been found, and the algorithm can be stopped.

In Yamamoto (1993) the algorithm of van den Elzen and Talman (1991) is explained using an expanding set which is a family of subsets $\Sigma(t)$ of Σ , where $\Sigma(t) = (1 - t)\{\nu\} + t\Sigma$ for $t \in [0, 1]$. By defining a series of games $\Gamma(t)$ with $t \in [0, 1]$ and with the space of action profiles defined above, by construction ν is a Nash equilibrium of the game $\Gamma(0)$. Moreover for increasing values of t , a path of Nash equilibria of the games $\Gamma(t)$ is generated that perfectly matches with the path in the procedure of van den Elzen and Talman. For a general interpretation of path-following algorithms as projections on appropriately defined expanding sets, see Herings (2002).

In our homotopy interpretation, it suffices to define the correspondence $H : [0, 1] \times \Sigma \rightarrow \Sigma$ by

$$H(t, \sigma) = \prod_{i \in N} \beta_{\Sigma^i(t)}^i(\sigma),$$

where $\beta_{\Sigma^i(t)}^i(\sigma) = \arg\max_{\bar{\sigma}^i \in \Sigma^i(t)} u^i(\sigma^{-i}, \bar{\sigma}^i)$ is the best response correspondence of player i in game Γ with restricted set of mixed strategies $\Sigma^i(t)$. The set of strategies of the resulting

homotopy path coincides with the set of strategies generated by the Van den Elzen-Talman algorithm. Convergence of the Van den Elzen-Talman algorithm for a nondegenerate game now follows easily from Theorem 2.2.

The formulation as a homotopy exemplifies the differences between the Lemke-Howson algorithm and the Van den Elzen-Talman one. Apart from the general starting point in Van den Elzen-Talman, the main difference is that for the starting and intermediate problems, Lemke-Howson leaves the strategy set of one player unrestricted, whereas Van den Elzen-Talman constrain both strategy sets. The formulation also clarifies that the Lemke-Howson problem can be formulated with a general starting point.

In van den Elzen and Talman (1999) it is shown that given any prior and a generic bimatrix game, the paths generated by the linear tracing procedure developed in Harsanyi (1975) and the pivoting procedure outlined above are the same. The tracing procedure is the key ingredient in the equilibrium selection method developed in Harsanyi and Selten (1988). We will explain the linear tracing procedure in more detail in Section 7.

6 n -Player Games: General Overview

The problem of finding a Nash equilibrium in a game with more than two players is equivalent to solving a non-linear complementarity problem. Therefore, linear programming techniques are in general not directly applicable. An exception is the class of polymatrix games, studied in Govindan and Wilson (2004). Polymatrix games are games where for each player the payoffs are additive over the bilateral interactions with the other players. A polymatrix game is therefore a linear combination of bimatrix games. Because of the linearity of the payoff functions, polymatrix games can be relatively quickly solved using a variant of the Lemke-Howson algorithm. By iteratively solving approximating polymatrix games (where in each iteration the polymatrix approximation used depends on the current mixed action profile) Govindan and Wilson exploit the linearity to increase the speed of their algorithm presented in Govindan and Wilson (2003).

For general n -person games, algorithms to compute equilibria rely either on simplicial subdivisions for approximating fixed points or numerical methods to solve systems of higher-degree polynomial inequalities. Methods for the computation of equilibria for noncooperative games with more than two players are surveyed in McKelvey and McLennan (1996).

Seminal papers on the computation of Nash equilibria in n -person games are the ones by Rosenmüller (1971) and Wilson (1971). Both extend the Lemke-Howson procedure to nondegenerate n -person games by formulating the problem of finding a Nash equilibrium as a nonlinear complementarity problem.

The Lemke-Howson procedure starts with fixing an action for one player and the best response against it for the other player. This implies solving a 1-person game for its equilibrium. This starting point determines an almost-complementary node for the 2-person game

from which a path towards an equilibrium originates.

The procedures of Rosenmüller (1971) and Wilson (1971) start with fixing an action for $(n - 1)$ players and compute the best response against it for the remaining player. The point found determines an almost-complementary node for the 2-person game that results after unfixing one of the players from its action from which a path towards an equilibrium of this 2-person game originates. This equilibrium determines the starting node for the next step in the procedure. The procedures of Rosenmüller and Wilson construct in this way an equilibrium of an n -person game by successively constructing equilibria of k -person games, $1 \leq k \leq n$.

The problem of the procedure is that the paths are no longer linear. To turn this procedure into a numerical algorithm, a method is needed to follow the non-linear path. Here, there are two options, one relying on simplicial approximations, explained in detail in Section 7, the other on differentiability, see Section 8.

Although the methods of Rosenmüller and Wilson are not directly suitable for computational purposes, their common result, the existence of a non-linear path leading to an equilibrium, was a very important step towards an implementable algorithm as developed by Garcia, Lemke, and Lüthi (1973). Later, a more efficient algorithm was proposed in van der Laan, Talman, and van der Heijden (1987). A problem of these algorithms is that they calculate only an approximation of a sample Nash equilibrium, and do not take the game-theoretic underpinning of the calculated equilibrium into consideration.

Another generalization of the Lemke-Howson algorithm to n -person games is provided by Govindan and Wilson (2003). They indicate that one of the implications of the structure theorem of Kohlberg and Mertens (1986) is that, above each generic ray emanating from the true game (represented as a point in a Euclidean space), the graph of the equilibrium correspondence is a 1-dimensional manifold. Moreover, at sufficient distance from the true game there is a unique equilibrium. Therefore, starting from a sufficiently distant game along any generic ray, one can traverse the line segment to the true game, tracing the 1-dimensional manifold of equilibria along the way, to find an equilibrium of the true game at the terminus. Wilson and Govindan (2003) propose to trace the manifold using a global Newton method. For bimatrix games the algorithm jumps linearly from one boundary to another, and corresponds to the Lemke-Howson algorithm.

Homotopy-based algorithms to solve for all Nash equilibria are developed in Kostreva and Kinard (1991) and Herings and Peeters (2005). Both methods are based on the application of numerical techniques to obtain all the solutions to a system of polynomial equations. Kostreva and Kinard (1991) focuses on solving for polynomial optimization problems in general and bimatrix games in particular. Herings and Peeters (2005) concentrates on solving n -person noncooperative games for all its Nash equilibria. The proposed method is shown to be globally convergent for an open set of games with full Lebesgue measure.

McLennan (2005) shows that the number of Nash equilibria in normal form games of

modest size is huge on average. For instance, the mean number of Nash equilibria in a game with 4 agents, each having 6 strategies, is estimated to be 2.037. This number increases rapidly in the number of players and the number of strategies. To compute all Nash equilibria of a game of moderate size is therefore not feasible within reasonable time limits. It is therefore important to have methods that compute a Nash equilibrium with a sound game-theoretic underpinning.

Yamamoto (1993) presents a homotopy to compute a proper Nash equilibrium. Yamamoto makes use of an expanding set to define a homotopy whose homotopy-path connects the centroid of the simplex to a solution of a stationary point problem that is a Nash equilibrium of the original game. For generic games, all Nash equilibria are proper. The ability to compute a proper Nash equilibrium does therefore not solve the problem equilibrium selection.

Two methods of equilibrium selection stand out in the game-theoretic literature. The first one is the equilibrium selection theory developed in Harsanyi and Selten (1988), which applies to general n -person games. The Herings-Van den Elzen and Herings-Peeters algorithms that are presented in Sections 7 and 8, respectively, compute the Nash equilibrium that is selected by the linear tracing procedure—an important construct in the selection theory of Harsanyi and Selten (1988). The second one is the (logit) quantal response equilibrium of McKelvey and Palfrey (1995). The quantal response equilibrium is often applied by experimentalists because of its good prediction of behavior. It incorporates probabilities of making mistakes into the Nash equilibrium concepts. Quantal response equilibrium can be turned into a theory of equilibrium selection by selecting the equilibrium that is obtained in the limit when mistake probabilities go to zero. We present this method in Section 9. A drawback of this selection theory is that the equilibrium selected is not invariant with respect to the chosen representation of the preferences by utilities.

Another equilibrium selection method that has a bounded rationality interpretation is the t -solution (Rosenthal, 1989). It relaxes the prescription that best responses are played with probability one, by replacing it with a probabilistic choice in which better responses are played with probabilities not lower than worse responses. The level of randomness of players' behavior is parameterized by a single parameter t that can be interpreted as a measure of the importance of the game. Voorneveld (2006) showed that for every game the set of t -solutions constitutes a homotopy path starting at $t = 0$ and terminating in a Nash equilibrium at $t = 1$. This path can be followed using the techniques presented in the next two sections. Moreover, Voorneveld (2006) shows that for increasing values of t , the players eliminate successively higher levels of never-best replies and that eventually the only actions that are chosen with positive probability are rationalizable.

7 n -Player Games: The Herings-Van den Elzen Algorithm

This section presents an algorithm that computes the Nash equilibrium selected by the linear tracing procedure of Harsanyi (1975). Harsanyi's linear tracing procedure starts from a prior, reflecting the ideas of any player about the action used by the other player, and next players react optimally thereupon. Then the players observe that their expectations are not met and adjust them towards observed behavior and again react thereupon. By simultaneously and gradually adjusting expectations, and reacting optimally against these revised expectations, eventually an equilibrium is reached.

Consider some n -person game Γ and some prior $p \in \Sigma$ and denote, for every $t \in [0, 1]$, Γ^t as the game $\langle N, \{S^i\}_{i \in N}, \{v^i(t)\}_{i \in N} \rangle$, where the payoff function $v^i(t) : \Sigma \rightarrow \mathbb{R}$ of player i is defined by

$$v^i(t, \sigma) = tu^i(\sigma) + (1 - t)u^i(p^{-i}, \sigma^i). \quad (8)$$

The game Γ^0 corresponds to a trivial game, where all players believe that their opponents play with probability 1 according to the prior belief. The game Γ^1 coincides with the original game Γ . A best response against a strategy combination $\sigma \in \Sigma$ in the game Γ^t corresponds to a best response against the probability distribution $t[\sigma] + (1 - t)[p]$ on S in the game Γ . The latter probability distribution does in general not belong to Σ , since it may be correlated when there are more than two players.

The set of all Nash equilibria related to the games Γ^t , $t \in [0, 1]$, is denoted by

$$\mathcal{L}(\Gamma, p) = \left\{ (t, \sigma) \in [0, 1] \times \Sigma \mid \sigma \in \text{NE}(\Gamma^t) \right\}. \quad (9)$$

Once the number of pure strategies of each player is fixed, a noncooperative game is completely determined by the utility functions u , which can be represented by a vector in $\mathbb{R}^{m^*n}$. The standard topology and measure on $\mathbb{R}^{m^*n}$ therefore induce a topology and a measure on games. The set of such games is denoted $\mathcal{G}$.

THEOREM 7.1 *There exists a path in $\mathcal{L}(\Gamma, p)$ connecting a best response against the prior to a Nash equilibrium of the game Γ . For an open set of games $\Gamma \in \mathcal{G}$ and priors $p \in \Sigma$ with full Lebesgue measure, this path is unique.*

For a proof of the first statement, see Schanuel, Simon and Zame (1991) and Herings (2000), for a proof of the latter statement, see Harsanyi (1975) and Herings and Peeters (2001). For a better understanding of the result, we define the homotopy correspondence $H : [0, 1] \times \Sigma \rightarrow \Sigma$ by

$$H(t, \sigma) = \prod_{i \in N} \beta^i(t; \sigma),$$

where as before $\beta^i(t; \sigma)$ is the set of best responses of player i in game Γ^t against strategy profile σ . The proof of the first statement is now an immediate application of Theorem 2.2.

The second statement needs that games are nondegenerate in a similar sense as before. It can be proved that this is the case for generic games. We will return to this issue in Section 8.

The linear tracing procedure links a Nash equilibrium of the game Γ^0 to a Nash equilibrium of Γ^1 . That is, the linear tracing procedure traces the generically unique path in $\mathcal{L}(\Gamma, p)$. The interpretation of the linear tracing procedure is that players gradually adjust their beliefs about the behavior of their opponents by gradually putting less weight to the initial beliefs, the prior.

Admissible subsets B^* can be used to decompose $\mathcal{L}(\Gamma, p)$ in subsets $\mathcal{L}(\Gamma, p, B^*)$, where a set $\mathcal{L}(\Gamma, p, B^*)$ contains those elements of $\mathcal{L}(\Gamma, p)$ where only strategies in B^* are played with positive probability. The set $\Sigma(B^*) = \{\sigma \in \Sigma \mid \forall s_j^i \notin B^*, \sigma_j^i = 0\}$ is the strategy set consistent with B^* . The set $\Lambda(B^*) = \{\lambda \in \mathbb{R}_+^{m^*} \mid \forall s_j^i \in B^*, \lambda_j^i = 0\}$ is the set of Lagrange multipliers associated to B^* . A point $(t, \sigma) \in [0, 1] \times \Sigma(B^*)$ belongs to $\mathcal{L}(\Gamma, p, B^*)$ if and only if there exists $\lambda \in \Lambda(B^*)$ and $\mu \in \mathbb{R}^n$ such that

$$v^i(t, \sigma^{-j}, s_j^i) + \lambda_j^i - \mu^i = 0 \quad (s_j^i \in S^*), \quad (10)$$

We denote by v the function with components $v^i(t, \sigma^{-j}, s_j^i)$ for $s_j^i \in S$. The next step is to turn Theorem 7.1 and Equations (10) into an implementable algorithm.

Because we are dealing with a non-linear system of equations, we have to study approximations of Nash equilibria. Our ultimate aim is to compute a ε -Nash equilibrium of Γ with ε an arbitrarily chosen positive number.

DEFINITION 7.2 Let (Γ, p) be given. For $\varepsilon \geq 0$, a mixed strategy combination $\sigma \in \Sigma$ is called an ε -Nash equilibrium of Γ if, for every $i \in N$, $\sigma_k^i > 0$ implies $u^i(\sigma^{-i}, s_k^i) \geq \max_{s_j^i \in S^i} u^i(\sigma^{-i}, s_j^i) - \varepsilon$.

In an ε -Nash equilibrium the loss in payoffs of using a suboptimal strategy is at most ε . Though the suboptimal strategy itself might be far away from an optimal strategy, the loss in payoff is small, which makes sense from a game-theoretic standpoint.

In Herings and van den Elzen (2002) a simplicial algorithm is used to generate a piecewise linear path that approximates $\mathcal{L}(\Gamma, p)$. For every B^* , they define a piecewise linear approximation of the Equations (10), whose solution approximates $\mathcal{L}(\Gamma, p, B^*)$. Next it is shown that the piecewise linear solutions paths for varying B^* can be nicely fitted together, and result in a piecewise linear path from a solution to the starting problem to an ε -Nash equilibrium of the terminal problem.

For $k \in \mathbb{N}$, a k -dimensional simplex or k -simplex ς in $\mathbb{R}^m$ is defined as the convex hull of $k + 1$ affinely independent points $x^1, \dots, x^{k+1}$ of $\mathbb{R}^m$. We write $\varsigma = \varsigma(x^1, \dots, x^{k+1})$ and call $x^1, \dots, x^{k+1}$ the vertices of ς . A $(k - 1)$ -simplex that is the convex hull of k vertices of $\varsigma(x^1, \dots, x^{k+1})$ is said to be a facet of ς . The facet $\tau(x^1, \dots, x^{j-1}, x^{j+1}, \dots, x^{k+1})$ is called the facet of ς opposite to the vertex x^j . For a non-negative integer k' less than or equal to k , a

k' -simplex that is the convex hull of $k' + 1$ vertices of ς is said to be a k' -face of ς .

A finite collection $\mathcal{S}$ of k -simplices is a triangulation of a k -dimensional convex subset T of $\mathbb{R}^m$ if (1) T is the union of all simplices in $\mathcal{S}$ and (2) the intersection of any two simplices in $\mathcal{S}$ is either empty or a common face of both. If $\mathcal{S}$ is a triangulation of T , and a facet τ of $\varsigma^1 \in \mathcal{S}$ is a subset of the relative boundary of T , then there is no $\varsigma^2 \in \mathcal{S}$ such that $\varsigma^2 \neq \varsigma^1$ and τ is a facet of ς^2 . If τ is not a subset of the relative boundary of T , then there is exactly one $\varsigma^2 \in \mathcal{S}$ such that $\varsigma^2 \neq \varsigma^1$ and τ is also a facet of ς^2 . The mesh size of a triangulation $\mathcal{S}$ of T is defined by $\text{mesh}(\mathcal{S}) = \max\{\|x - y\|_\infty \mid x, y \in \varsigma, \varsigma \in \mathcal{S}\}$.

It is well-known that full-dimensional affine parts of the relative boundary of a set are triangulated by the facets of the simplices in a triangulation. More precisely, let $\mathcal{S}$ be a triangulation of a convex k -dimensional subset T of $\mathbb{R}^m$, and let the $(k-1)$ -dimensional subset $\bar{T}$ of the relative boundary of T be such that $\bar{T}$ is equal to the affine hull of $\bar{T}$ intersected with T . Then the collection $\bar{\mathcal{S}}$ given by $\bar{\mathcal{S}} = \{\tau \in \bar{T} \mid \exists \varsigma \in \mathcal{S}, \tau \text{ is a facet of } \varsigma\}$ is a triangulation of $\bar{T}$, see Todd (1976), Theorem 2.3. For instance, the set $\{0\} \times \Sigma$ is triangulated by the facets of the simplices in a triangulation of $[0, 1] \times \Sigma$.

An example of a triangulation of $[0, 1] \times \Sigma$ is illustrated in Figure 4 for the case we have two players each having two pure strategies. All 3-simplices in $\text{co}(\bar{x}^3, \bar{x}^5, \bar{x}^6, \bar{x}^{12}, \bar{x}^{14}, \bar{x}^{15})$ are depicted in Figure 4. For later purposes we give all 3-simplices in the triangulation of Figure 4

Figure 4: A triangulation of $[0, 1] \times \Sigma$. The vertex $\bar{x}^1 = (0, (1, 0), (1, 0))$ refers to the strategy vector at which both players play their first pure strategy. Similarly, $\bar{x}^3$, $\bar{x}^7$, and $\bar{x}^9$ correspond to $(0, (0, 1), (1, 0))$, $(0, (1, 0), (0, 1))$, and $(0, (0, 1), (0, 1))$, respectively.

in Table 1. In Table 1 only the 3-simplices in $[0, \frac{1}{2}] \times \Sigma$ are given. The ones in $[\frac{1}{2}, 1] \times \Sigma$ follow by means of a translation. The position in the table is related to the position of a simplex in

the triangulation.

$\text{co}(\{\bar{x}^4, \bar{x}^7, \bar{x}^5, \bar{x}^{14}\})$	$\text{co}(\{\bar{x}^8, \bar{x}^7, \bar{x}^5, \bar{x}^{14}\})$	$\text{co}(\{\bar{x}^8, \bar{x}^9, \bar{x}^5, \bar{x}^{14}\})$	$\text{co}(\{\bar{x}^6, \bar{x}^9, \bar{x}^5, \bar{x}^{14}\})$
$\text{co}(\{\bar{x}^4, \bar{x}^7, \bar{x}^{16}, \bar{x}^{14}\})$	$\text{co}(\{\bar{x}^8, \bar{x}^7, \bar{x}^{16}, \bar{x}^{14}\})$	$\text{co}(\{\bar{x}^8, \bar{x}^9, \bar{x}^{18}, \bar{x}^{14}\})$	$\text{co}(\{\bar{x}^6, \bar{x}^9, \bar{x}^{18}, \bar{x}^{14}\})$
$\text{co}(\{\bar{x}^4, \bar{x}^{13}, \bar{x}^{16}, \bar{x}^{14}\})$	$\text{co}(\{\bar{x}^8, \bar{x}^{17}, \bar{x}^{16}, \bar{x}^{14}\})$	$\text{co}(\{\bar{x}^8, \bar{x}^{17}, \bar{x}^{18}, \bar{x}^{14}\})$	$\text{co}(\{\bar{x}^6, \bar{x}^{15}, \bar{x}^{18}, \bar{x}^{14}\})$
$\text{co}(\{\bar{x}^4, \bar{x}^1, \bar{x}^5, \bar{x}^{14}\})$	$\text{co}(\{\bar{x}^2, \bar{x}^1, \bar{x}^5, \bar{x}^{14}\})$	$\text{co}(\{\bar{x}^2, \bar{x}^3, \bar{x}^5, \bar{x}^{14}\})$	$\text{co}(\{\bar{x}^6, \bar{x}^3, \bar{x}^5, \bar{x}^{14}\})$
$\text{co}(\{\bar{x}^4, \bar{x}^1, \bar{x}^{10}, \bar{x}^{14}\})$	$\text{co}(\{\bar{x}^2, \bar{x}^1, \bar{x}^{10}, \bar{x}^{14}\})$	$\text{co}(\{\bar{x}^2, \bar{x}^3, \bar{x}^{12}, \bar{x}^{14}\})$	$\text{co}(\{\bar{x}^6, \bar{x}^3, \bar{x}^{12}, \bar{x}^{14}\})$
$\text{co}(\{\bar{x}^4, \bar{x}^{13}, \bar{x}^{10}, \bar{x}^{14}\})$	$\text{co}(\{\bar{x}^2, \bar{x}^{11}, \bar{x}^{10}, \bar{x}^{14}\})$	$\text{co}(\{\bar{x}^2, \bar{x}^{11}, \bar{x}^{12}, \bar{x}^{14}\})$	$\text{co}(\{\bar{x}^6, \bar{x}^{15}, \bar{x}^{12}, \bar{x}^{14}\})$

Table 1: All full-dimensional simplices in $[0, \frac{1}{2}] \times S$.

A function $\bar{v} : [0, 1] \times \Sigma \rightarrow \mathbb{R}^{m^*}$ is called a piecewise linear approximation of v with respect to $\mathcal{S}$ if for each vertex x^k of any $\varsigma(x^1, \dots, x^{m^*-n+2}) \in \mathcal{S}$, $\bar{v}(x^k) = v(x^k)$ and $\bar{v}$ is affine on each simplex of $\mathcal{S}$. Hence, if $x \in \varsigma(x^1, \dots, x^{m^*-n+2})$, so $x = \sum_{k=1}^{m^*-n+2} \alpha_k x^k$, $\alpha_k \geq 0$, $k = 1, \dots, m^* - n + 2$, $\sum_{k=1}^{m^*-n+2} \alpha_k = 1$, then $\bar{v}(x) = \sum_{k=1}^{m^*-n+2} \alpha_k \bar{v}(x^k)$.

Let an admissible subset B^* and a triangulation $\mathcal{S}$ of $[0, 1] \times \Sigma$ be given. We denote by $\mathcal{S}(B^*)$ the collection of b -faces of simplices in Σ , where $b = |B^*| + 1 - n$, that are contained in $[0, 1] \times \Sigma(B^*)$ and $|B^*|$ denotes the cardinality of B^* . By repeated application of the result that claims that the relative boundary of a set is triangulated by the facets of a triangulation, it follows that $\mathcal{S}(B^*)$ is a triangulation of $[0, 1] \times \Sigma(B^*)$.

Let a simplex $\varsigma(x^1, \dots, x^{b+1}) \in \mathcal{S}(B^*)$ be given. Consider solutions $(\alpha, \lambda, \mu) \in \mathbb{R}_+^{b+1} \times \Lambda(B^*) \times \mathbb{R}^n$ of the following system of equations, a piecewise linear approximation of the system (10),

$$\begin{aligned} \sum_{k=1}^{b+1} \alpha_k &= 1, \\ \sum_{k=1}^{b+1} \alpha_k \bar{v}^i(x^k) + \sum_{s_j^i \notin B^*} \lambda_j^i e_j^i - \mu^i e^i &= 0^i, \quad i \in N. \end{aligned} \tag{11}$$

Such solutions are called admissible. An admissible solution (α, λ, μ) corresponds to an approximate Nash equilibrium σ of Γ^t . Indeed, $(t, \sigma^\top)^\top = \sum_{k=1}^{b+1} \alpha_k x^k$, strategies in B^* are best replies for the payoff function $\bar{v}$, λ_j^i is the payoff gap between strategy s_j^i and a best reply for player i , and μ^i is the payoff for player i according to $\bar{v}$ when he uses a best reply. Since $\varsigma \subset [0, 1] \times \Sigma(B^*)$, strategies that are not a best reply are played with probability zero. A solution to (11) is said to be degenerate if at least two of the variables α_k , $k = 1, \dots, b + 1$, and λ_j^i , $s_j^i \notin B^*$, are equal to zero.

The Herings-Van den Elzen algorithm generates by means of lexicographic pivoting techniques a piecewise linear path of approximate Nash equilibria in $[0, 1] \times \Sigma$ joining $\{0\} \times \Sigma$ to $\{1\} \times \Sigma$. The path is such that every (t, σ) on it corresponds to an admissible B^* , a simplex $\varsigma \in \mathcal{S}(B^*)$, and a solution (α, λ, μ) . The Herings-Van den Elzen algorithm specifies in a unique way how to move from one simplex to another.

For given (B^*, ς) , (11) corresponds to a linear system with $m^* + 1$ equations and $m^* + 2$ variables. If we rule out degeneracies, then a non-empty solution set is a 1-dimensional

compact line segment. The end-points of the line segment are either approximate Nash equilibria for Γ^0 or Γ^1 , or yield solutions for a new $(\bar{B}^*, \bar{\varsigma})$. Indeed, with degeneracies ruled out, at an end-point either $\alpha_k = 0$ for exactly one k or $\lambda_j^i = 0$ for exactly one $s_j^i \notin B^*$. In the first case, the end-point belongs to the facet τ of ς opposite to the vertex x^k . If τ belongs to the relative interior of $[0, 1] \times \Sigma(B^*)$, then there is a unique simplex $\bar{\varsigma} \in \mathcal{S}(B^*)$ such that $\bar{\varsigma} \neq \varsigma$, and τ is a facet of $\bar{\varsigma}$. The algorithm continues by generating a line-segment of solutions in $\bar{\varsigma}$. If τ belongs to the relative boundary of $[0, 1] \times \Sigma(B^*)$, then the end-point is either an approximate Nash equilibrium for Γ^0 , or an approximate Nash equilibrium for Γ^1 , or $\tau \in \mathcal{S}(\bar{B}^*)$, with $\bar{B}^*$ a uniquely determined subset of B^* having one element less, and the algorithm continues with a line-segment of solutions in τ . If $\lambda_j^i = 0$, then also strategy s_j^i is a best reply. The algorithm continues with a line-segment of solutions in $\bar{\varsigma}$, where $\bar{\varsigma}$ is the unique simplex in $\mathcal{S}(B^* \cup \{s_j^i\})$ having ς as a facet.

In game theory degeneracy is not always a non-generic phenomenon. For instance, for a normal form representation of a game in extensive form, degeneracy is the rule rather than the exception, even if the payoffs in the extensive form game are randomly chosen. But also in other normal form games, representing certain economic situations, degeneracy can easily occur, simply because payoffs are not randomly chosen but reflect some structure that is present in the economic model. Degeneracy can be dealt with by exploiting lexicographic pivoting techniques. We explain next how lexicographic pivoting techniques can be used to extend the ideas of the previous paragraph to handle degenerate games.

For an admissible B^* and a facet $\tau(x^1, \dots, x^b)$ of a simplex in $\mathcal{S}(B^*)$, the $(m^*+1) \times (m^*+1)$ -matrix $A_{B^*, \tau}$ is defined by

$$A_{B^*, \tau} = \begin{bmatrix} 1 & \cdots & 1 & 0 & \cdots & \cdots & \cdots & \cdots & 0 \\ & & & E^1 & & 0 & -e^1 & & 0 \\ \bar{v}(x^1) & \cdots & \bar{v}(x^b) & 0 & \ddots & 0 & 0 & \ddots & 0 \\ & & & 0 & 0 & E^n & 0 & 0 & -e^n \end{bmatrix},$$

where, for $i \in N$, $E^i = [e_j^i]_{s_j^i \notin B^*}$ and the zeros indicate submatrices with zeros of appropriate dimension. The matrix $A_{B^*, \tau}$ corresponds to the coefficients in (11) when a facet τ of a simplex ς is considered. Suppose $A_{B^*, \tau}^{-1}$ exists. From $A_{B^*, \tau} A_{B^*, \tau}^{-1} = I^{m^*+1}$ (the (m^*+1) -dimensional identity matrix), it follows that the first column of $A_{B^*, \tau}^{-1}$ corresponds to an admissible solution to (11) for any $\varsigma \in \mathcal{S}(B^*)$ being the convex hull of τ and some vertex $x^{b+1} \in [0, 1] \times \Sigma(B^*)$, whenever the first m^*+1-n components of this column are non-negative. No restrictions are imposed on the last n rows of $A_{B^*, \tau}^{-1}$. In a nondegenerate solution the first m^*+1-n components are all strictly positive, since $\alpha_{b+1} = 0$ extends the solution for the facet τ to the simplex ς .

A row vector $x \in \mathbb{R}^{m^*+1}$ is lexicographically positive if it is not equal to the vector of zeroes and its first non-zero entry is positive. The matrix $A_{B^*, \tau}^{-1}$ is said to be semi-lexicopositive if

each of the first $m^* + 1 - n$ rows is lexicographically positive.

DEFINITION 7.3 Let $(\Gamma, p, \mathcal{S})$ and an admissible B^* be given. A facet τ of a simplex in $\mathcal{S}(B^*)$ is B^* -complete if $A_{B^*, \tau}^{-1}$ exists and is semi-lexicopositive.

Given a linear system of equations as in (11), one can pivot in a uniquely determined new column, either of the type $(1, \bar{v}(x^{b+1}))$ or of the type $(0, 0, e_j^i, 0)$. In the nondegenerate case such a pivot step determines in a unique way a column out of the first $m^* + 1 - n$ to be replaced. In degenerate cases it is possible that the leaving column is not uniquely determined. A semi-lexicographic pivot step is a pivot step where the leaving column is selected in such a way that the inverse of the resulting matrix $A_{B^*, \tau}$ is semi-lexicopositive. Herings and Van den Elzen (2002) show that a semi-lexicographic pivot step determines in all cases a unique column out of the first $m^* + 1 - n$ to be replaced.

Theorem 7.4 describes all possible cases that may occur if a B^* -complete facet τ is given and a semi-lexicographic pivot step with a vector $(1, \bar{v}(x^{b+1}))$ is made, where x^{b+1} is a vertex of a simplex having τ as a facet opposite to it.

THEOREM 7.4 Let $(\Gamma, p, \mathcal{S})$ and a B^* -complete facet τ of a simplex $\varsigma \in \mathcal{S}(B^*)$ be given. Then exactly one of the following cases holds:

1. ς has exactly one other B^* -complete facet $\bar{\tau}$,
2. ς is $\bar{B}^*$ -complete for precisely one admissible $\bar{B}^*$.

Theorem 7.5 describes all possible cases that may occur if a B^* -complete facet τ is given that is also a simplex belonging to $\mathcal{S}(\bar{B}^*)$, where $B^* = \bar{B}^* \cup \{s_j^i\}$, and a semi-lexicographic pivot step with a vector $(0, 0, e_j^i, 0)$ is made.

THEOREM 7.5 Let $(\Gamma, p, \mathcal{S})$ and a B^* -complete facet τ that belongs to $\mathcal{S}(\bar{B}^*)$ for some admissible $\bar{B}^*$ be given. Then exactly one of the following cases holds:

1. τ is $\hat{B}^*$ -complete for precisely one admissible $\hat{B}^*$ with $\hat{B}^* \neq B^*$,
2. precisely one facet v of τ is $\bar{B}^*$ -complete.

The consideration of B^* -complete facets determines a unique starting point for the algorithm. The admissible subset B^{*0} is defined by the set of strategies s_j^i , where j is the largest integer such that s_j^i is a best reply to the prior p for player i . Notice that $|B^{*0}| = n$. It can be shown that the facet (vertex) $\tau = \{0\} \times \Sigma(B^{*0})$ is B^{*0} -complete and that there is no other B^* -complete facet τ in $\{0\} \times \Sigma$. Even in degenerate cases, the semi-lexicographic rules single out the unique B^{*0} -complete facet $\{0\} \times \Sigma(B^{*0})$, which serves as a unique starting point of the algorithm.

Notice that we are now in a position that makes the Lemke-Howson door-in door-out principle applicable. There is a unique starting point in $\{0\} \times \Sigma$, the B^{*0} -complete facet $\{0\} \times \Sigma(B^{*0})$. For any B^* -complete facet τ of a simplex in $\mathcal{S}(B^*)$, Theorems 7.4 and 7.5 determine another complete facet in a unique way. The finiteness of the number of simplices guarantees that at some stage a facet that is a subset of $\{1\} \times \Sigma$ is reached. Such a facet determines an approximate Nash equilibrium.

The formal steps of the Van den Elzen-Herings algorithm are as follows.

ALGORITHM 7.6 Let $(\Gamma, p, \mathcal{S})$ be given.

Step 0. Let $b = 1$ and $r = 1$. Let $B^* = B^{*0}$, $\tau^1 = \{0\} \times \Sigma(B^*)$, and let x^2 be the unique vertex of the 1-simplex of $\mathcal{S}(B^*)$ containing τ^1 as the facet opposite to it.

Step 1. Let ς be equal to the convex hull of τ^r and $\{x^{b+1}\}$. Make a semi-lexicographic pivot step with $(1, \bar{v}(x^{b+1}))$ into the system of equations (11) corresponding to A_{B^*, τ^r} , yielding a unique column k' of A_{B^*, τ^r} which has to be replaced. If $k' \in \{b+1, \dots, m^*+1-n\}$, then go to Step 3 with $s_{j'}^{i'}$, the pure strategy corresponding to column k' . Otherwise, go to Step 2.

Step 2. Increase the value of r by 1 and let τ^r be the facet of ς opposite $x^{k'}$. If $\tau^r \subset \{1\} \times \Sigma$, then the algorithm terminates with an approximate Nash equilibrium σ^* of Γ^1 induced by the solution of (11) corresponding to A_{B^*, τ^r} . If $\tau^r \in \mathcal{S}(\bar{B}^*)$ for some admissible $\bar{B}^*$, then go to Step 4. Otherwise, there is exactly one b -simplex $\bar{\varsigma}$ of $\mathcal{S}(B^*)$ such that $\bar{\varsigma} \neq \varsigma$ and τ^r is a facet of $\bar{\varsigma}$. Go to Step 1 with x^{b+1} as the unique vertex of $\bar{\varsigma}$ opposite τ^r .

Step 3. Let the admissible $\bar{B}^*$ be defined by $\bar{B}^* = B^* \cup \{s_{j'}^{i'}\}$. There is a unique simplex $\bar{\varsigma}$ of $\mathcal{S}(\bar{B}^*)$ having ς as a facet. Increase the value of both b and r by 1 and go to Step 1 with x^{b+1} as the unique vertex of $\bar{\varsigma}$ opposite ς , $B^* = \bar{B}^*$, and $\tau^r = \varsigma$.

Step 4. Let ς be equal to τ^r . Make a semi-lexicographic pivot step with $(0, 0, e_{j^*}^{i^*}, 0)$ into the system of equations (11) corresponding to A_{B^*, τ^r} , where $s_{j^*}^{i^*}$ is such that $\bar{B}^* \cup \{s_{j^*}^{i^*}\} = B^*$. This yields a unique column k' of A_{B^*, τ^r} which has to be replaced. If $k' \in \{b+1, \dots, m^*+1-n\}$, then decrease the value of both b and r by 1 and go to Step 3 with $s_{j'}^{i'}$, the pure strategy corresponding to column k' and $B^* = \bar{B}^*$. Otherwise, decrease the value of b by 1 and go to Step 2 with $B^* = \bar{B}^*$.

Theorem 7.4 corresponds to the semi-lexicographic pivot step made in Step 1 of Algorithm 7.6. Case 1 of Theorem 7.4 occurs if one goes from Step 1 of Algorithm 7.6 to Step 2, and Case 2 if one goes from Step 1 to Step 3. Theorem 7.5 corresponds to the semi-lexicographic pivot step performed in Step 4 of Algorithm 7.6. Case 1 of Theorem 7.5 happens if one goes from Step 4 of Algorithm 7.6 to Step 3, and Case 2 if one goes from Step 4 to Step 2. The algorithm

terminates after a finite number of steps, after having generated a B^* -complete facet τ being a subset of $\{1\} \times \Sigma(B^*)$.

The algorithm generates a piecewise linear approximation of the homotopy path generated by the linear tracing procedure. For every $\varepsilon > 0$, there is $\delta > 0$ such that a triangulation of $[0, 1] \times \Sigma$ with mesh size smaller than δ makes the piecewise linear approximation generated by the Herings-Van den Elzen algorithm within ε -Hausdorff distance from the homotopy path generated by the tracing procedure and hence terminates with a ε -Nash equilibrium of Γ . If the latter path is unique, and a sequence of triangulations with decreasing mesh size is taken, then the generated piecewise linear approximations converge to the tracing procedure path in the Hausdorff topology.

In Figure 5 the algorithm is illustrated for the game considered before. The prior p is taken to be $((\frac{1}{2}, \frac{1}{2}), (\frac{3}{4}, \frac{1}{4}))$. Since p^2 is part of the Nash equilibrium, both pure strategies of player 1 are best responses, and this game is degenerate in the sense of Van den Elzen-Talman. We have chosen this prior to illustrate that lexicographic pivoting techniques overcome degeneracies. Since the game is a bimatrix game, we would in general recommend to use either the Lemke-Howson or the Van den Elzen-Talman algorithm, rather than the Herings-Van den Elzen algorithm.

Figure 5: The Herings-Van den Elzen algorithm in action.

In the example the algorithm generates 13 facets before terminating with an approximate Nash equilibrium. The 2-dimensional facets generated are shaded in Figure 5. The path generated by the algorithm is illustrated by the heavily drawn line going from τ^1 to τ^{13} . The dotted line represents $\mathcal{L}(\Gamma, p)$, which for this example consists of one component. It can be decomposed into six sets $\mathcal{L}(\Gamma, p, B^*)$.

8 n -Person Games: The Herings-Peeters Algorithm

Games with n players possess a nice piecewise differentiable structure that can be exploited for computational purposes. This is exactly what the Herings-Peeters algorithm does. This algorithm exploits the manifold structure of the sets $\mathcal{L}(\Gamma, p, B^*)$.

THEOREM 8.1 *For an open set of games $\Gamma \in \mathcal{G}$ and priors $p \in \Sigma$ with full Lebesgue measure, for all admissible subsets B^* of S^* , the set $\mathcal{L}(\Gamma, p, B^*)$ is a compact 1-dimensional C^∞ manifold with boundary. Moreover, (t, σ) is a boundary point of $\mathcal{L}(\Gamma, p, B^*)$ if and only if either $\sigma_j^i = 0$ for exactly one $s_j^i \in B^*$, or exactly one $s_j^i \notin B^*$ is a best response to σ^{-i} , or $t = 0$, or $t = 1$.*

A compact 1-dimensional C^∞ manifold with boundary consists of finitely many arcs and loops. Theorem 8.1 therefore guarantees that generically $\mathcal{L}(\Gamma, p, B^*)$ has a simple mathematical structure, excluding bifurcations, spirals, higher dimensional parts, etc. Theorem 8.1 is the non-linear equivalent of the compact line segments of the Lemke-Howson, Van den Elzen-Talman and Van den Elzen-Herings algorithms. One difference now is that loops are possible and, moreover, $\mathcal{L}(\Gamma, p, B^*)$ may contain more than one arc or loop.

To prove Theorem 8.1, the theory of regular constraint sets as presented in Jongen, Jonker and Twilt (1983) can be applied, see also Herings (1997) for a first application of this theory to economics. For some $r \geq 1$ a subset M of $\mathbb{R}^m$ is called a k -dimensional C^r manifold with generalized boundary (MGB), if for every $\bar{x} \in M$ there exists a C^r diffeomorphism $\varphi : U \rightarrow V$, where U is an open subset of $\mathbb{R}^m$ containing $\bar{x}$ and V is open in $\mathbb{R}^m$, and some integer $\ell(\bar{x}) \geq 0$, such that $\varphi(\bar{x}) = 0$ and $\varphi(U \cap M)$ equals $\{y \in V \mid y_l = 0, l = 1, \dots, m - k, \text{ and } y_l \geq 0, l = m - k + 1, \dots, m - k + \ell(\bar{x})\}$. If for every element $\bar{x}$ of an MGB M it holds that $\ell(\bar{x}) \leq 1$, then M is called a manifold with boundary and the set of elements $\bar{x}$ for which $\ell(\bar{x}) = 1$ is an $(k - 1)$ -dimensional manifold, called the boundary of M .

Let K^1 and K^2 be two finite index sets and let g_k for all $k \in K^1$ and h_k for all $k \in K^2$, be C^r functions defined on some open subset X of $\mathbb{R}^m$. We define

$$M[g, h] = \{x \in X \mid g_k(x) = 0, \forall k \in K^1, \text{ and } h_k(x) \geq 0, \forall k \in K^2\}.$$

For $x \in X$ we define the set $K^0(x) = \{k \in K^2 \mid h_k(x) = 0\}$. If for every $\bar{x} \in M[g, h]$ it holds that

$$\{\partial_x g_k(\bar{x}) \mid k \in K^1\} \cup \{\partial_x h_k(\bar{x}) \mid k \in K^0(\bar{x})\} \quad (12)$$

is a set of independent vectors, then $M[g, h]$ is called a C^r regular constraint set (RCS). In Jongen, Jonker and Twilt (1983) it is shown that every C^r RCS is an $(m - |K^1|)$ -dimensional C^r MGB with $\ell(\bar{x}) = |K^0(\bar{x})|$ for every $\bar{x} \in M[g, h]$. If a set is a k -dimensional manifold with generalized boundary, then the neighborhood of a point $\bar{x}$ belonging to that set looks, in a well-defined sense, like $\mathbb{R}^{k-\ell(\bar{x})} \times \mathbb{R}_+^{\ell(\bar{x})}$.

We can therefore show $\mathcal{L}(\Gamma, p, B^*)$ to be a C^∞ 1-dimensional manifold with boundary, if we can represent it as a C^∞ regular constraint set with index sets K^1 and K^2 , functions g_k for $k \in K^1$ holding with equality, and functions h_k for $k \in K^2$ holding with weak inequality, where $|K^1| = m - 1$. Moreover, we have to show that $|J^0(\bar{x})| = 1$ for every $\bar{x} \in M[g, h]$.

A point (t, σ) belongs to $\mathcal{L}(\Gamma, p, B^*)$ if and only if there exists $\lambda \in \mathbb{R}^{m^*}$ and $\mu \in \mathbb{R}^n$ such that

$$\begin{aligned}
v^i(t, \sigma^{-i}, s_j^i) + \lambda_j^i - \mu^i &= 0 & (s_j^i \in S^*) \\
\sigma_j^i &= 0 & (s_j^i \notin B^*) \\
\lambda_j^i &= 0 & (s_j^i \in B^*) \\
\sum_{s_j^i \in S^i} \sigma_j^i - 1 &= 0 & (i \in N) \\
\sigma_j^i &\geq 0 & (s_j^i \in B^*) \\
\lambda_j^i &\geq 0 & (s_j^i \notin B^*) \\
t &\geq 0 \\
-t + 1 &\geq 0.
\end{aligned} \tag{13}$$

The set K^1 corresponding to (13) has cardinality $2m^* + n$. This is indeed one less than the number of free variables, which is equal to $2m^* + n + 1$. Unfortunately, (13) is not a regular constraint set in general. The set corresponding to (12) may not consist of independent vectors. Moreover, it is easy to construct examples such that $|J^0(\bar{x})| \geq 2$ for some $\bar{x} \in M[g, h]$. Fortunately, such examples can be shown to be exceptional cases. To prove that, a result known as the transversality theorem comes to rescue, see for instance Mas-Colell (1985). Let C^1 manifolds X , Y , and Z , Z being a subset of Y , an element $\bar{x}$ of X , and a function $f \in C^1(X, Y)$ be given. The function f is said to intersect Z transversally at $\bar{x} \in X$, denoted by $f \bar{\cap} Z$ at $\bar{x}$, if

$$f(\bar{x}) \notin Z, \text{ or } f(\bar{x}) \in Z \text{ and } T_{f(\bar{x})}Z + \partial f(\bar{x})(T_{\bar{x}}X) = T_{f(\bar{x})}Y,$$

where $T_{\bar{x}}X$ denotes the tangent space of X at $\bar{x}$. For regular constraint sets, the tangent space of X at $\bar{x}$ is easily computed as

$$T_{\bar{x}}X = \{x \in \mathbb{R}^m \mid \partial g(\bar{x})(x) = 0\},$$

where 0 is an $(m - k)$ -dimensional vector of zeroes with k being the dimension of the tangent space. The function f is said to intersect Z transversally if $f \bar{\cap} Z$ at every $x \in X$. Transversality of f implies that the inverse image of f has a particularly nice structure.

THEOREM 8.2 *For $k^1, k^2, k^3 \in \{0\} \cup \mathbb{N}$, for $r \in \mathbb{N} \cup \{\infty\}$, let a k^1 -dimensional C^1 manifold X , a k^2 -dimensional C^1 manifold Y , and a k^3 -dimensional C^1 manifold Z , Z being a subset of Y , be given, and let the function $f \in C^1(X, Y)$ be such that $f \bar{\cap} Z$. If $k^1 - k^2 + k^3 < 0$, then*

$f^{-1}(Z) = \emptyset$. If $f \in C^r(X, Y)$, $f \pitchfork Z$, and $k^1 - k^2 + k^3 \geq 0$, then $f^{-1}(Z)$ is a $(k^1 - k^2 + k^3)$ -dimensional C^r manifold.

We are now in a position to state the transversality theorem.

THEOREM 8.3 For $k^1, k^2, k^3 \in \{0\} \cup \mathbb{N}$, for $r \in \mathbb{N} \cup \{\infty\}$, let a k^1 -dimensional C^r manifold X , a k^2 -dimensional C^r manifold Y , a k^3 -dimensional C^r manifold Z being a subset of Y , a C^r manifold Ω , and a function $f \in C^r(X \times \Omega, Y)$ be given, with $r \geq \max(\{1, k^1 - k^2 + k^3\})$. For every $\omega \in \Omega$, define a function $f^\omega \in C^r(X, Y)$ by $f^\omega(x) = f(x, \omega)$, $\forall x \in X$. Then $f \pitchfork Z$ implies $f^\omega \pitchfork Z$, except for a subset of Ω having Lebesgue measure zero in Ω .

We apply Theorem 8.3 in the following way. We take as the set of parameters Ω the set of games and priors. Then we specify three types of regular constraint sets. The first type consists of all equalities of (13), the second type of all equalities plus one inequality of (13), and the third type of all equalities plus two inequalities of (13). Thus we obtain three types of functions f , with variables $t, \sigma, \lambda, \mu, \omega$, that can be shown to satisfy $f \pitchfork \{0\}$, and as a consequence $f^\omega \pitchfork \{0\}$ for almost every $\omega \in \Omega$. Counting the number of equations and unknowns, using Theorem 8.2, it follows that a function f^ω of the third type, with two inequalities, can only be transversal if $f^{\omega^{-1}}\{0\} = \emptyset$. From this we derive the conclusion that, generically, it is impossible that two inequality constraints are binding at the same time. The transversality of the first two types of functions leads to the conclusion that, generically, (13) is a regular constraint set. This concludes the basic steps in the proof of Theorem 8.1.

Using the representation of $\mathcal{L}(\Gamma, p, B^*)$ as a regular constraint set, we also find that a point (t, σ) is a boundary point of $\mathcal{L}(\Gamma, p, B^*)$ if and only if either $\sigma_j^i = 0$ for exactly one $s_j^i \in B^*$ or exactly one $s_j^i \notin B^*$ is a best response to σ^{-i} , or $t = 0$, or $t = 1$. Moreover, two sets $\mathcal{L}(\Gamma, p, B^*)$ and $\mathcal{L}(\Gamma, p, \bar{B}^*)$ can only have boundary points in common, since otherwise (13) would have a solution with two inequality constraints binding at the same time. When there is a common boundary point, there is an optimal action s_j^i that is played with zero probability, with the two admissible subsets only differing in this action: $(B^* \cup \bar{B}^*) \setminus (B^* \cap \bar{B}^*) = \{s_j^i\}$. Moreover, in each point at most two sets can meet, as we would again obtain a case where (13) has two inequality constraints binding at the same time. We see that all subsets $\mathcal{L}(\Gamma, p, B^*)$ of $\mathcal{L}(\Gamma, p)$ are nicely connected. What we have obtained is a full non-linear analogue of the arguments underlying the Lemke-Howson algorithm.

THEOREM 8.4 For an open set of games $\Gamma \in \mathcal{G}$ and priors $p \in \Sigma$ with full Lebesgue measure, $\mathcal{L}(\Gamma, p)$ is a compact one-dimensional piecewise C^∞ manifold with boundary. All boundary points of $\mathcal{L}(\Gamma, p)$ are in $\{0, 1\} \times \Sigma$. The boundary point in $\{0\} \times \Sigma$ is unique.

For almost every Γ and p , the set $\mathcal{L}(\Gamma, p)$ consists of a finite number of paths and loops.

Although it is not necessarily the case that these paths and loops are smooth, the number of non-differentiabilities is finite at most. All paths in $\mathcal{L}(\Gamma, p)$ start and end in $\{0, 1\} \times \Sigma$. Each such path consists of a finite sequence of smooth arcs of the sets $\mathcal{L}(\Gamma, p, B^*)$. A loop in $\mathcal{L}(\Gamma, p)$ consists either of a finite sequence (at least two) of differentiable arcs in the sets $\mathcal{L}(\Gamma, p, B^*)$ or is a loop of one set $\mathcal{L}(\Gamma, p, B^*)$.

Since $\mathcal{L}(\Gamma, p)$ consists of finitely many 1-dimensional C^∞ manifolds with boundary, it is possible to use standardized software based on predictor-corrector methods, to approximately follow each manifold. We refer the interested reader to Allgower and Georg (1990) for a detailed exposition on predictor-corrector methods. Although this is certainly feasible, it is attractive to avoid switching from one system of equations to the next, and to formulate one, differentiable, system of equations whose solution corresponds to $\mathcal{L}(\Gamma, p)$.

Define for $\alpha \in \mathbb{R}^{m^*}$, differentiable and monotone (and strictly monotone on the domain where they are non-zero) functions $\sigma_j^i(\alpha)$ and $\lambda_j^i(\alpha)$ such that $\sigma_j^i(\alpha) \geq 0$, $\lambda_j^i(\alpha) \geq 0$, and $\sigma_j^i(\alpha) \cdot \lambda_j^i(\alpha) = 0$. Next consider the system

$$\begin{aligned} v^i(t, \sigma^{-i}(\alpha), s_j^i) + \lambda_j^i(\alpha) - \mu^i &= 0 & (s_j^i \in S^*) \\ \sum_{s_j^i \in S^i} \sigma_j^i(\alpha) - 1 &= 0 & (i \in N) \\ t &\geq 0 \\ -t + 1 &\geq 0. \end{aligned} \tag{14}$$

For each point (t, α, μ) satisfying the (in)equalities (14), the point $(t, \sigma(\alpha), \lambda(\alpha), \mu)$ satisfies the (in)equalities (13) with $B^* = \{s_j^i \in S^* \mid \sigma_j^i(\alpha) > 0\}$.

Let $\mathcal{H} : [0, 1] \times \mathbb{R}^{m^*} \times \mathbb{R}^n \rightarrow \mathbb{R}^{m^*} \times \mathbb{R}^n$ be the continuously differentiable homotopy function defined by

$$\mathcal{H}(t, \alpha, \mu) = \begin{pmatrix} v^i(t, \sigma^{-i}(\alpha), s_j^i) + \lambda_j^i(\alpha) - \mu^i & (s_j^i \in S^i, i \in N) \\ \sum_{s_j^i \in S^i} \sigma_j^i(\alpha) - 1 & (i \in N) \end{pmatrix}. \tag{15}$$

The zeros of this homotopy function describe the linear tracing procedure: $(t, \alpha, \mu) \in \mathcal{H}^{-1}(\{0\})$ if and only if $(t, \sigma(\alpha)) \in \mathcal{L}(\Gamma, p)$. Starting at the unique point $(0, \alpha^0, \mu^0) \in \mathcal{H}^{-1}(\{0\})$ at $t = 0$ and following the path described by the zeros of $\mathcal{H}$, we end up in a point $(1, \tilde{\alpha}, \tilde{\mu}) \in \mathcal{H}^{-1}(\{0\})$. This point generates the Nash equilibrium $\sigma(\tilde{\alpha})$ of Γ selected by the Harsanyi-Selten theory.

THEOREM 8.5 *For an open set of games $\Gamma \in \mathcal{G}$ and priors $p \in \Sigma$ with full Lebesgue measure, $\mathcal{H}^{-1}(\{0\})$ is a compact 1-dimensional C^1 manifold with boundary. All boundary points of $\mathcal{H}^{-1}(\{0\})$ are in $\{0, 1\} \times \mathbb{R}^{m^*} \times \mathbb{R}^n$. The boundary point in $\{0\} \times \mathbb{R}^{m^*} \times \mathbb{R}^n$ is unique.*

The set $\mathcal{H}^{-1}(\{0\})$ consists of finitely many differentiable arcs and loops. All arcs start and end in $\{0, 1\} \times \mathbb{R}^{m^*} \times \mathbb{R}^n$. Loops have no points in common with $\{0, 1\} \times \mathbb{R}^{m^*} \times \mathbb{R}^n$. There is exactly one arc that starts in $\{0\} \times \mathbb{R}^{m^*} \times \mathbb{R}^n$ and that ends in $\{1\} \times \mathbb{R}^{m^*} \times \mathbb{R}^n$ with a point $(1, \tilde{\alpha}, \tilde{\mu})$ that generates the Nash equilibrium selected by the tracing procedure. This arc is a

transformation of the feasible path of the tracing procedure. All other arcs start and end in $\{1\} \times \mathbb{R}^{m^*} \times \mathbb{R}^n$ and connect two points inducing Nash equilibria of Γ .

The structure of $\mathcal{H}^{-1}(\{0\})$ is even simpler than the one of $\mathcal{L}(\Gamma, p)$. Not only, like for $\mathcal{L}(\Gamma, p)$, are complications like bifurcations, spirals, higher dimensional solutions sets, diverging behavior, etc. excluded. The arcs and loops in $\mathcal{H}^{-1}(\{0\})$ are differentiable everywhere. It is the transformation of variables that smoothes out the kinks. As a direct consequence, it is possible to calculate the derivative at each point of the feasible path, which makes it possible to follow the path by means of differentiable as opposed to simplicial methods. In Herings and Peeters (2001) we report numerical results using the software-package HOMPACK, a FORTRAN77 program. This program implements three predictor-corrector methods, see Watson, Billups, and Morgan (1987) for an introduction to HOMPACK.

For the game in Figure 1 and as prior the mixed action combination $p = ((\frac{1}{2}, \frac{1}{2}), (\frac{1}{2}, \frac{1}{2}))$, the corresponding set $\mathcal{L}(\Gamma, p)$ is given by

$$\begin{aligned} \mathcal{L}(\Gamma, p) = & \{ (t, (s_2^1, s_1^2)) \mid t \in [0, \frac{1}{2}] \} \cup \{ (\frac{1}{2}, ((r, 1-r), s_1^2)) \mid r \in [0, \frac{5}{6}] \} \\ & \cup \{ (t, ((\frac{3t+1}{6t}, \frac{3t-1}{6t}), (\frac{2t+1}{4t}, \frac{2t-1}{4t}))) \mid t \in [\frac{1}{2}, 1] \}. \end{aligned} \quad (16)$$

At $t = 0$, for player 1 and player 2 the actions s_2^1 and s_1^2 respectively are the best responses against the prior. The first part of the path of $\mathcal{L}(\Gamma, p)$ is therefore determined by $\mathcal{L}(\Gamma, p, B^*)$ for $B^* = \{s_2^1, s_1^2\}$. As long as $t \leq \frac{1}{2}$ the path stays at (s_2^1, s_1^2) . At $t = \frac{1}{2}$ action s_1^1 of player 1 becomes optimal and has to be added to the admissible set B^* such that the path continues in $\mathcal{L}(\Gamma, p, B^*)$ for $B^* = \{s_1^1, s_2^1, s_1^2\}$. In this new component, at $t = \frac{1}{2}$ the weight on player 1's action s_1^1 is increased until either it gets weight 1 or until player 2's unused action s_2^2 becomes optimal. This latter happens when the weight on action s_1^1 reaches the value $\frac{5}{6}$. The final part of the path is determined by the component $\mathcal{L}(\Gamma, p, B^*)$ with B^* including all actions.

The first plot of Figure 6 shows the values of t , σ and λ along the homotopy path that is generated by the linear tracing procedure as a function of pathlength. To make the plot more clear, only informative values of the variables are depicted. In particular this means that for each pure strategy s_j^i , either σ_j^i or $-\lambda_j^i$ is plotted, depending on which one is non-zero. We see that at two values of τ kinks in the paths occur. These values of τ correspond to the two points where we had to adapt the admissible subset B^* in the description above.

The second plot shows the values of t and α as a function of pathlength and therefore corresponds to the feasible path of the tracing procedure after application of the transformation $\sigma_j^i(\alpha) = [\max\{0, \alpha_j^i\}]^2$ and $\lambda_j^i(\alpha) = [\max\{0, -\alpha_j^i\}]^2$. From the plot it is seen that the transformation indeed causes the kinks in the path to disappear. At the points where previously kinks occurred, now the derivatives of all variables with respect to pathlength are equal to zero, except the variable that passes zero at that point. It is precisely the transformation of variables that does the trick.

Figure 6: The homotopy path before and after transformation.

9 n -Person Games: The McKelvey-Palfrey Algorithm

Quantal response equilibria as introduced by McKelvey and Palfrey (1995) are statistical versions of Nash equilibria where each player's payoff is subject to random error. For a given specification of the error structure, a quantal response equilibrium is a mixed action combination that is consistent with optimizing behavior subject to that error structure. One possible interpretation of this is that players make calculation errors according to some random process when calculating their expected payoffs.

Player i 's payoff when playing pure strategy s_j^i against a mixed strategy combination σ^{-i} is subject to an error and is given by

$$\hat{u}^i(\sigma^{-i}, s_j^i) = u^i(\sigma^{-i}, s_j^i) + \varepsilon_j^i. \quad (17)$$

Player i 's error vector ε^i is distributed according to a joint distribution with density function $\varphi^i(\varepsilon^i)$. The marginal distribution of φ^i exists for each ε_j^i and $\mathbb{E}(\varepsilon^i) = 0$. For any possible opponents' mixed action combination σ^{-i} , the s_j^i -response set R_j^i is defined as the set of error vectors that make strategy s_j^i the best response, so

$$R_j^i(\sigma^{-i}) = \{ \varepsilon^i \mid \hat{u}^i(\sigma^{-i}, s_j^i) \geq \hat{u}^i(\sigma^{-i}, s_\ell^i) \text{ for all } s_\ell^i \in S^i \}. \quad (18)$$

The probability of choosing pure action s_j^i is then given by

$$\pi_j^i(\sigma^{-i}) = \int_{R_j^i(\sigma^{-i})} \varphi^i(\varepsilon^i) d\varepsilon^i. \quad (19)$$

The function π^i is called the statistical reaction function, or the quantal response function and satisfies the feature that better actions are more likely to be chosen than worse actions.

A *quantal response equilibrium* (QRE) is a mixed action combination σ that is consistent with the error structure:

$$\sigma_j^i = \pi_j^i(\sigma^{-i}) \quad \text{for all } s_j^i \in S^i \text{ and all } i \in N. \quad (20)$$

THEOREM 9.1 (MCKELVEY AND PALFREY, 1995) *For any game Γ and for any density function φ , a QRE exists.*

For any parameter $\lambda \geq 0$, the *logistic quantal response function* is defined by

$$\pi_j^i(\sigma^{-i}) = \frac{\exp(\lambda u^i(\sigma^{-i}, s_j^i))}{\sum_{s_\ell^i \in S^i} \exp(\lambda u^i(\sigma^{-i}, s_\ell^i))} \quad (21)$$

and is obtained when φ^i corresponds to the extreme value (or log-Weibull) distribution. Therefore, if each player uses a logistic quantal response function, the corresponding QRE, called *logit equilibrium*, requires that

$$\sigma_j^i = \frac{\exp(\lambda u^i(\sigma^{-i}, s_j^i))}{\sum_{s_\ell^i \in S^i} \exp(\lambda u^i(\sigma^{-i}, s_\ell^i))} \quad \text{for all } s_j^i \in S^i \text{ and all } i \in N. \quad (22)$$

The parameter λ is inversely related to the error level. When $\lambda = 0$, the choice of the players is completely determined by the errors which induces all players to play all their pure actions with equal probability. When λ approaches infinity, the influence of the error disappears.

The *logit equilibrium correspondence* $\mathcal{L} : \mathbb{R}_+ \rightarrow \Sigma$ is given by

$$\mathcal{L}(\lambda) = \left\{ \sigma \in \Sigma \mid \sigma_j^i = \frac{\exp(\lambda u^i(\sigma^{-i}, s_j^i))}{\sum_{s_\ell^i \in S^i} \exp(\lambda u^i(\sigma^{-i}, s_\ell^i))} \quad (s_j^i \in S^i, i \in N) \right\}. \quad (23)$$

THEOREM 9.2 (MCKELVEY AND PALFREY, 1995) *When λ approaches infinity, the set of logit equilibria converges to a subset of the set of Nash equilibria. For an open set of games $\Gamma \in \mathcal{G}$ with full Lebesgue measure, the graph of $\mathcal{L}$ contains a unique path of logit equilibria that starts at the centroid for $\lambda = 0$ and converges to a unique Nash equilibrium as λ goes to infinity.*

The latter unique Nash equilibrium is called the *limiting logit equilibrium* of the game Γ . This induces a unique selection from the set of Nash equilibria by ‘tracing’ the graph of the logit equilibrium correspondence beginning at the centroid of the strategy simplex and continuing for increasing values of λ .

Define the correspondence $\tilde{\mathcal{L}}$ from $[0, 1]$ to Σ by

$$\tilde{\mathcal{L}}(t) = \begin{cases} \mathcal{L}(\frac{t}{1-t}) & \text{if } t \in [0, 1) \\ \text{NE}(\Gamma) & \text{if } t = 1 \end{cases} \quad (24)$$

THEOREM 9.3 (HERINGS, 2002) *For all games $\Gamma \in \mathcal{G}$ the graph of $\tilde{\mathcal{L}}$ contains a component containing an element at $t = 0$ and at $t = 1$.*

The (homotopy) path of logit equilibria can be traced from the centroid at $t = 0$ towards the limiting logit equilibrium at $t = 1$ using one of the two methods described in the previous two sections.

10 Extensive Form and Dynamic Games

Since extensive form games can be represented as a normal-form game, in principle it is possible to apply any of the algorithms of the previous sections. There are two caveats. First, since the normal-form representation of an extensive form game is non-generic, it becomes crucial to handle degeneracies appropriately. Second, from a computational point of view, it might be much more attractive to work with the extensive form representation. Here, we will restrict ourselves to a brief overview of this important stream of the literature.

Wilson (1972) applied the Lemke-Howson algorithm to a two-person extensive form game with perfect recall. By making use of the extensive form directly, the developed method never deals with the entire linear complementary problem, and leads to significant savings in storage.

In the same journal issue, von Stengel (1996) and Koller, Megiddo and von Stengel (1996) propose algorithms to solve two-person extensive form games with perfect recall: the first for the class of such games with zero-sum payoff structure, the latter for general payoffs. Perfect recall implies that players never forget what they once knew or what they did earlier. For games with perfect recall mixed strategies are essentially equivalent to behavioral strategies. This allows for studying the sequence form, rather than the normal form, of the extensive form. The advantage of the sequence form in comparison to the normal form is that the size of the sequence form is linear and not exponentially in the size of the game tree. For the resulting problem the algorithm of Lemke-Howson is applied.

Later von Stengel, van den Elzen and Talman (2002) present yet another algorithm for solving two-person extensive form games with perfect recall. Just like the earlier contributions by von Stengel *c.s.* the sequence form is used rather than the normal form. In this contribution, instead of applying the Lemke-Howson algorithm, the van den Elzen-Talman algorithm is used. The advantages of that method are the potential to find multiple equilibria and to find normal form perfect equilibria.

McKelvey and Palfrey (1998) extend the logit quantal response equilibrium defined for normal form games to the logit agent quantal response equilibrium (AQRE) for extensive form games. In an AQRE, at each information set players choose better actions with higher probabilities than worse actions. They show that limit points of the logit AQRE yield a refine-

ment of sequential equilibria for any finite extensive form game, but are not logically related to other refinement criteria (such as the intuitive criterion or trembling hand perfection).

For stochastic games, homotopy based algorithms have been proposed by Filar and Raghavan (1984), Nowak and Raghavan (1993), and Raghavan and Syed (2002). All papers concentrate on two-person games with the single-controller property (i.e. only one player has partial control over the state transitions): the first for discounted games with zero-sum payoff structure, the second for discounted games with general payoff structure, the third for undiscounted games. In all papers a stationary equilibrium is defined as a solution to a linear complementary problem, where linearity is obtained as a result of the single-controller property. Subsequently, the Lemke-Howson algorithm is applied to solve for a stationary equilibrium.

Herings and Peeters (2004) extend the linear tracing procedure to a related procedure that is suitable for general finite discounted stochastic games. The extension is shown to possess the same geometric properties as the linear tracing procedure possesses for normal form games. By numerically following the exact homotopy path, which is shown to be generically unique, an approximation of a stationary equilibrium results. Since this method allows for an arbitrary starting point, it has the potential to find multiple equilibria.

11 Conclusions

The paper presents an overview of homotopy algorithms as applied to non-cooperative game theory. Advantages of homotopy algorithms include their numerical stability, their ability to locate multiple solutions, and the insight they provide in the properties of solutions. Homotopy algorithms can be implemented easily with the aid of existing, professionally programmed, software.

From an algorithmic point of view, it is useful to distinguish bi-matrix games from games with more than two players. The reason is that the problem of finding a Nash equilibrium in a bi-matrix games is equivalent to solving a linear complementarity problem, whereas the general problem is equivalent to solving a non-linear complementarity problem.

We present two algorithms in detail that are suitable to compute Nash equilibria for bi-matrix games, the Lemke-Howson algorithm and the Van den Elzen-Talman algorithm. We argue that both algorithms are best understood as following a solution path generated by a homotopy. We also explain how the door-in door-out principle of Lemke-Howson applies to both algorithms.

For general n -person games, it is usually not possible to compute an equilibrium exactly, which calls for different methods. We present the two main ideas in detail, one using simplicial methods, the other predictor-corrector methods. The Herings-Van den Elzen algorithm relies on the simplicial approach, the Herings-Peeters algorithm on the predictor-corrector idea. Both algorithms converge to an approximate Nash equilibrium for general n -person games.

For both methods we also illustrate how they can be understood as following a solution path generated by a homotopy. We also explain how the door-in door-out principle of Lemke-Howson applies to both algorithms. The third method we propose to find a Nash-equilibrium of a general n -person game is related to the quantal-response equilibrium of McKelvey and Palfrey.

Since the number of Nash equilibria of an arbitrary game tends to be enormous, we emphasize the importance of computing a Nash equilibrium with a good game-theoretical underpinning. The algorithms of Van den Elzen-Talman, Herings-Van den Elzen, and Herings-Peeters are all related to the equilibrium selection methods of Harsanyi and Selten. The McKelvey-Palfrey algorithm has an interesting behavioral interpretation.

We conclude with a brief overview of how extensions of these methods have been applied to games in extensive form and to dynamic games. We think that this is a still underexplored research area and we expect more exciting work to be done in the future.

References

1. Allgower, E.L., and K. Georg (1990), *Numerical Continuation Methods: An Introduction*, Springer-Verlag, Berlin, Heidelberg.
2. Browder, F.E. (1960), "On Continuity of Fixed Points under Deformations of Continuous Mappings," *Summa Brasiliensis Mathematicae*, 4, 183-191.
3. Eaves, B.C., and K. Schmedders (1999), "General Equilibrium Models and Homotopy Methods," *Journal of Economic Dynamics & Control*, 23, 1249-1279.
4. Elzen, A.H. van den, 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.
5. Elzen, A.H. van den, and A.J.J. Talman (1999), "An Algorithmic Approach toward the Tracing Procedure for Bi-matrix Games," *Games and Economic Behavior*, 28, 130-145.
6. Filar, J.A., and T.E.S. Raghavan (1984), "A Matrix Game Solution of a Single Controller Stochastic Game," *Mathematics of Operations Research*, 9, 356-362.
7. Garcia, C.B., C.E. Lemke, and H.J. Lüthi (1973), "Simplicial Approximation of an Equilibrium Point of Noncooperative N -Person Games," in T.C. Hu and S.M. Robinson (eds.), *Mathematical Programming*, Academic Press, New York, pp. 227-260.
8. Garcia, C.B., and W.I. Zangwill (1981), *Pathways to Solutions, Fixed Points, and Equilibria*, Prentice Hall, Englewood Cliffs.

9. Geanakoplos, J.D. (2003), "Nash and Walras Equilibrium Via Brouwer," *Economic Theory*, 21, 585-603.
10. Govindan, S., and R. Wilson (2003), "A Global Newton Method to Compute Nash Equilibria," *Journal of Economic Theory*, 110, 65-86.
11. Govindan, S., and R. Wilson (2004), "Computing Nash Equilibria by Iterated Polymatrix Approximation," *Journal of Economic Dynamics & Control*, 28, 1229-1241.
12. 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.
13. Harsanyi, J.C., and R. Selten (1988), *A General Theory of Equilibrium Selection in Games*, MIT Press, Cambridge.
14. Herings, P.J.J. (1997), "A Globally and Universally Stable Price Adjustment Process," *Journal of Mathematical Economics*, 27, 163-193.
15. Herings, P.J.J. (2000), "Two Simple Proofs of the Feasibility of the Linear Tracing Procedure," *Economic Theory*, 15, 485-490.
16. Herings, P.J.J. (2002), "Universally Converging Adjustment Processes—A Unifying Approach," *Journal of Mathematical Economics*, 38, 341-370.
17. Herings, P.J.J., and A.H. van den Elzen (2002), "Computation of the Nash Equilibrium Selected by the Tracing Procedure in N -Person Games," *Games and Economic Behavior*, 38, 89-117.
18. Herings, P.J.J., and R. Peeters (2001), "A Differentiable Homotopy to Compute Nash Equilibria of n -Person Games," *Economic Theory*, 18, 159-185.
19. Herings, P.J.J., and R. Peeters (2004), "Stationary Equilibria in Stochastic Games: Structure, Selection and Computation," *Journal of Economic Theory*, 118, 32-60.
20. Herings, P.J.J., and R. Peeters (2005), "A Globally Convergent Algorithm to Compute all Nash Equilibria for n -Person Games," *Annals of Operations Research*, 137, 349-368.
21. Jongen, H.Th., P. Jonker, and F. Twilt (1983), *Nonlinear Optimization in $\mathbb{R}^n$, I. Morse Theory, Chebyshev Approximation*, Methoden und Verfahren der Mathematischen Physik, 29, Peter Lang, Frankfurt.
22. Judd, K.L. (1997), "Computational Economics and Economic Theory: Substitutes or Complements?," *Journal of Economic Dynamics and Control*, 21, 907-942.

23. Kohlberg, E., and J.F. Mertens (1986), "On the Strategic Stability of Equilibria," *Econometrica*, 54, 1003-1037.
24. Koller, D., N. Megiddo, and B. von Stengel (1996), "Efficient Computation of Equilibria for Extensive Two-Person Games," *Games and Economic Behavior*, 14, 247-259.
25. Kostreva, M.M., and L.A. Kinard (1991), "A Differential Homotopy Approach for Solving Polynomial Optimization Problems and Noncooperative Games," *Computers & Mathematics with Applications*, 21, 135-143.
26. Laan, G. van der, A.J.J. Talman, and L. van der Heijden (1987), "Simplicial Variable Dimension Algorithms for Solving the Nonlinear Complementarity Problem on a Product of Unit Simplices Using a General Labelling," *Mathematics of Operation Research*, 12, 377-397.
27. Lemke, C.E., and J.T. Howson Jr. (1964), "Equilibrium Points of Bimatrix Games," *SIAM Journal on Applied Mathematics*, 12, 413-423.
28. Mas-Colell, A. (1974), "A Note on a Theorem of F. Browder," *Mathematical Programming*, 6, 229-233.
29. Mas-Colell, A. (1985), *The Theory of General Economic Equilibrium, A Differentiable Approach*, Cambridge University Press, Cambridge.
30. McKelvey, R.D., and A. McLennan (1996), "Computation of Equilibria in Finite Games," In: *Handbook of Computational Economics, Volume I*, H.M. Amman, D.A. Kendrick, and J. Rust (eds.), Elsevier Science.
31. McKelvey, R.D., and T.R. Palfrey (1995), "Quantal Response Equilibria for Normal Form Games," *Games and Economic Behavior*, 10, 6-38.
32. McKelvey, R.D., and T.R. Palfrey (1998), "Quantal Response Equilibria for Extensive Form Games," *Experimental Economics*, 1, 9-41.
33. McLennan, A. (2005), "The Expected Number of Nash Equilibria of a Normal Form Game," *Econometrica*, 73, 141-174.
34. Nowak, A.S., and T.E.S. Raghavan (1993), "A Finite Step Algorithm via a Bimatrix Game to a Single Controller Nonzero-Sum Stochastic Game," *Mathematical Programming*, 59, 249-259.
35. Raghavan, T.E.S., and Z. Syed (2002), "Computing Stationary Nash Equilibria of Undiscounted Single-Controller Stochastic Games," *Mathematics of Operations Research*, 22, 384-400.

36. Rosenmüller, J. (1971), "On a Generalization of the Lemke-Howson Algorithm to Non-cooperative N -Person Games," *SIAM Journal on Applied Mathematics*, 21, 73-79.
37. Rosenthal, R.W. (1989), "A Bounded-Rationality Approach to the Study of Noncooperative Games," *International Journal of Game Theory*, 18, 273-292.
38. Schanuel, S.H., L.K. Simon, and W.R. Zame (1991), "The Algebraic Geometry of Games and the Tracing Procedure," in R. Selten (ed.), *Game Equilibrium Models II, Methods, Morals and Markets*, Springer-Verlag, Berlin, 9-43.
39. Shapley, L.S. (1974), "A Note on the Lemke-Howson Algorithm," *Mathematical Programming Study 1: Pivoting and Extensions*, 175-189.
40. Stengel, B. von (1996), "Efficient Computation of Behavior Strategies," *Games and Economic Behavior*, 14, 220-246.
41. Stengel, B. von (2002), "Computing Equilibria for Two-Person Games," In: *Handbook of Game Theory, Vol. 3, Chapter 45*, R.J. Aumann and S. Hart (eds.), North-Holland, Amsterdam, pp. 1723-1759.
42. Stengel B. von, 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.
43. Todd, M.J. (1976), *The Computation of Fixed Points and Applications*, Lecture Notes in Economics and Mathematical Systems, 124, Springer-Verlag, Berlin.
44. Voorneveld, M. (2006), "Probabilistic Choice in Games: Properties of Rosenthal's t -Solutions," *International Journal of Game Theory*, 34, 105-122.
45. Watson, L.T., S.C. Billups, and A.P. Morgan (1987), "HOMPACK: A Suite of Codes for Globally Convergent Homotopy Algorithms," *ACM Transactions on Mathematical Software*, 13, 281-310.
46. Wilson, R. (1971), "Computing Equilibria of N -Person Games," *SIAM Journal on Applied Mathematics*, 21, 80-87.
47. Wilson, R. (1972), "Computing Equilibria of Two-Person Games from the Extensive Form," *Management Science*, 18, 448-460.
48. Wilson, R. (1992), "Computing Simply Stable Equilibria," *Econometrica*, 60, 1039-1070.
49. Yamamoto, Y. (1993), "A Path-Following Procedure to Find a Proper Equilibrium of Finite Games," *International Journal of Game Theory*, 22, 249-259.