Stationary Equilibria in Stochastic Games: Structure , Selection, and Computation

Citation for published version (APA):

Herings, P. J. J., & Peeters, R. J. A. P. (2004). Stationary Equilibria in Stochastic Games: Structure , Selection, and Computation. *Journal of Economic Theory*, 118, 32-60.
<https://doi.org/10.1016/j.jet.2003.10.001>

Document status and date:

Published: 01/01/2004

DOI:

[10.1016/j.jet.2003.10.001](https://doi.org/10.1016/j.jet.2003.10.001)

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.

ELSEVIER

Available online at www.sciencedirect.com

SCIENCE @ DIRECT®

Journal of Economic Theory 118 (2004) 32–60

JOURNAL OF
**Economic
Theory**

<http://www.elsevier.com/locate/jet>

Stationary equilibria in stochastic games: structure, selection, and computation

P. Jean-Jacques Herings* and Ronald J.A.P. Peeters

Department of Economics, Maastricht University, P.O. Box 616, 6200 MD Maastricht, The Netherlands

Received 28 August 2002; final version received 8 October 2003

Abstract

This paper introduces an algorithm to compute stationary equilibria in stochastic games that is guaranteed to converge for almost all such games. Since in general the number of stationary equilibria is overwhelming, we pay attention to the issue of equilibrium selection. We do this by extending the linear tracing procedure to the class of stochastic games, called the stochastic tracing procedure. As a by-product of our results, we extend a recent result on the generic finiteness of stationary equilibria in stochastic games to oddness of equilibria.

© 2003 Elsevier Inc. All rights reserved.

JEL classification: C63; C72; C73

Keywords: Game theory; Stochastic games; Computation of equilibria; Linear tracing procedure

1. Introduction

Many economic situations of interest can be modelled as a stochastic game. Recent work of, for instance, Olley and Pakes [25], Pakes and Ericson [26], Pakes and McGuire [27], and Bergemann and Välimäki [1] is devoted to the application of stochastic games to problems emerging in the industrial organization literature. In our opinion, further progress in this research program can be made by developing methods to solve stochastic games numerically. Numerical solution methods allow researchers to go back and forth between the implications of economic theory and the characteristics of alternative datasets. See also [17,21] for an expression of the

*Corresponding author.

E-mail address: p.herings@algec.unimaas.nl (P.J.J. Herings).

important role of computational methods in the further development of economic theory.

The aim of this paper is to present an algorithm to compute a stationary equilibrium for an arbitrary finite n -person discounted stochastic game. Even though such a game may possess many non-stationary equilibria, there are good reasons to focus on stationary equilibria. Several motivations for analyzing stationary equilibria can be found in [20]. Stationary strategies prescribe the simplest form of behavior that is consistent with rationality, stationarity captures the notion that ‘bygones are bygones’ more completely than does the concept of subgame-perfect equilibrium, and it embodies the principle that ‘minor causes should have minor effects’, that is, only those aspects of the past that are ‘significant’ should have an appreciable influence on behavior. The pragmatic motivations they give are that in applied theory, the focus on stationary strategies allows for clean, unobstructed analysis of the influence of the state variables, that stationary strategies substantially reduce the number of parameters to be estimated in dynamic (econometric) models, and that stationary models can be simulated.

Stochastic games were introduced by Shapley [32]. He considered both finite and infinite horizon two-person zero-sum stochastic games with finite state space and finite action spaces. Shapley proved that such games have a value and that both players possess optimal stationary strategies with respect to the discounted payoff criterion.

Fink [6], Takahashi [35], and Sobel [33] extended Shapley’s model to general n -person stochastic games. For the model with a finite state space and finite action spaces they showed the existence of a stationary equilibrium. In [2,3,5,31], the problem of finding discounted stationary equilibria in a general n -person stochastic game is reduced to that of finding a global minimum in a nonlinear program with linear constraints. Solving this nonlinear program is equivalent to solving a certain nonlinear system for which it is known that the objective value in the global minimum is zero. But, as is noted in [2], the convergence of an optimization algorithm to the global optimum is not guaranteed. The same convergence problems occur in the work of Pakes and McGuire [27], who consider the case with continuous strategy choices for the players. Since their algorithm is essentially an iterated best reply algorithm, it cannot compute equilibria involving mixed strategies, which implies nonconvergence for the substantial class of stochastic games without equilibria in pure strategies.

For many normal form games there is a vast multiplicity of equilibria, see [24]. There is no reason to expect that the situation is different for the multiplicity of stationary equilibria in stochastic games. For this reason, we also find it essential to provide an algorithm that selects a particular stationary equilibrium. All computational methods mentioned so far do not make a systematic selection from the set of equilibria. If there are multiple equilibria, the computed equilibrium will typically depend on the starting point, which makes equilibrium selection essentially random.

In this paper we propose an algorithm that is shown to converge to a stationary equilibrium for a generic n -person stochastic game. The algorithm also deals with the

equilibrium selection problem in that it computes the stationary equilibrium selected by a natural extension of the linear tracing procedure of Harsanyi and Selten [10], which we call the stochastic tracing procedure.

In [10] the tracing procedure is defined for normal form games and for extensive form games with a perfect recall information structure. Algorithms to compute the Nash equilibrium selected by the tracing procedure in normal form games are proposed in [4,13,14]. For extensive form games, Harsanyi and Selten [10] first transform the game into one in standard form and subsequently define the tracing procedure for that class of games. Computation of the Nash equilibrium selected by the tracing procedure in extensive form games is the topic of von Stengel et al. [34], who invoke the sequence form to calculate such equilibria efficiently. Since expected utility in stationary strategies does not hold in stochastic games, it is not possible to transform a stochastic game into one in standard form. The way to extend the tracing procedure to the class of stochastic games is neither straightforward nor unique.

The algorithm belongs to the class of homotopy methods, a class well-known in economics thanks to the pioneering work of Scarf [30]. Apart from dealing with the equilibrium selection problem, another advantage of our approach is that we have been able to specify a differentiable homotopy. This makes it possible to apply standard path-following techniques that are available in professionally programmed software. This makes implementation on a computer an easy exercise. As a by-product of the proof that the algorithm converges for a generic stochastic game, we obtain the result that for a generic stochastic game the stochastic tracing procedure yields a path leading to a unique stationary equilibrium.

As another by-product of our convergence proof, we obtain an extension of a recent result of Haller and Lagunoff [8]. Their main result is that the set of stationary equilibria in a stochastic game is generically finite. A corollary to our main result is that this finite number of equilibria is odd.

An alternative to the approach taken in this paper, would be to extend the quantal-response method of McKelvey and Palfrey [22,23] to the class of stochastic games. This would be very interesting as well, since it would also give rise to a systematic method of equilibrium selection in stochastic games. If it is also possible to develop an everywhere differentiable homotopy that computes that equilibrium deserves to be studied in future work.

The paper has been organized as follows. In Section 2 we describe the model of the finite discounted stochastic game. In Section 3 the restriction to stationary strategies is made and the set of stationary equilibrium strategies is characterized. The definition of the stochastic tracing procedure is given in Section 4. In Section 5, the properties of the stochastic tracing procedure are studied. It is shown that for almost every stochastic game, the stochastic tracing procedure is formed by a finite union of arcs and loops. Using a well-chosen transformation of variables, the stochastic tracing procedure is described by the zeros of an everywhere differentiable homotopy function in Section 6. Section 7 discusses the implementation of the homotopy algorithm and provides some numerical results.

2. Finite discounted stochastic games

In this paper we study finite discounted stochastic games. Such a game is defined as follows.

Definition 2.1. A *finite discounted stochastic game* is an ordered sextuple

$$\Gamma = \langle N, \Omega, \{S_{\omega}^i\}_{(i,\omega) \in N \times \Omega}, \{u^i\}_{i \in N}, \pi, \delta \rangle,$$

where N , Ω and S_{ω}^i are finite non-empty sets, $S_{\omega} = \times_{i \in N} S_{\omega}^i$, u^i is a real-valued function on the set $H = \{(\omega, s_{\omega}) \mid \omega \in \Omega, s_{\omega} \in S_{\omega}\}$, where π is a map $\pi: H \rightarrow \Delta(\Omega)$ with $\Delta(\Omega)$ the family of probability distributions on the space Ω and δ is a discount factor.

The game parameters have the following meaning.

- $N = \{1, \dots, n\}$ is the player set.
- $\Omega = \{\omega_1, \dots, \omega_z\}$ is the state space.
- $S_{\omega}^i = \{s_{\omega 1}^1, \dots, s_{\omega m_i}^i\}$ is the action set of player $i \in N$ in state $\omega \in \Omega$.
- $u^i: H \rightarrow \mathbb{R}$ is the payoff function of player $i \in N$; if in state $\omega \in \Omega$ the players action choices are $s_{\omega} = (s_{\omega}^1, \dots, s_{\omega}^n) \in S_{\omega} = \times_{i \in N} S_{\omega}^i$, then player i gets an instantaneous payoff of $u^i(\omega, s_{\omega})$.
- $\pi: H \rightarrow \Delta(\Omega)$ is the transition map. For each $(\omega, s_{\omega}) \in H$, we can identify $\pi(\omega, s_{\omega})$ with the vector $(\pi(\omega_1 \mid \omega, s_{\omega}), \dots, \pi(\omega_z \mid \omega, s_{\omega}))$. Here $\pi(\bar{\omega} \mid \omega, s_{\omega})$ represents the probability that the system jumps to state $\bar{\omega}$ if in state ω the strategy-tuple s_{ω} is played. Hence, $\pi(\bar{\omega} \mid \omega, s_{\omega}) \geq 0$ and $\sum_{\bar{\omega} \in \Omega} \pi(\bar{\omega} \mid \omega, s_{\omega}) = 1$.
- $\delta \in (0, 1)$ is the discount factor and is used to discount future payoffs.

Such a stochastic game corresponds to a dynamic system which can be in different states and where at certain stages the players can influence the course of the play. We consider the infinite horizon model and the set of stages is assumed to be identical with the set $\mathbb{Z}_+ = \{0, 1, \dots\}$. Players know the game itself and that this knowledge is common knowledge among all the players. Moreover, the initial state ω^0 at stage $k = 0$ is common knowledge to the players.

The game proceeds as follows. All players i select at the initial state, simultaneously and independently of each other (possibly by a chance experiment) an action $s_{\omega^0}^i \in S_{\omega^0}^i$. Now two things happen, both depending on the current state ω^0 and the action choices $s_{\omega^0}^1, \dots, s_{\omega^0}^n$,

- (a) player i earns $u^i(\omega^0, s_{\omega^0})$,
- (b) the system jumps to the next state ω^1 according to the outcome of a chance experiment. The probability that the next state will be $\bar{\omega}$ equals $\pi(\bar{\omega} \mid \omega^0, s_{\omega^0})$.

Subsequently, prior to the next stage $k = 1$, all players are informed about the previous actions chosen by the players, and of the new state ω^1 . At stage $k = 1$, the above procedure is repeated, starting from the state ω^1 .

We assume that the game is of perfect recall, i.e., at each stage each player remembers all past actions chosen by all players and all past states that have occurred. Note that for finite stochastic games, each stage game resembles a normal form game. However, contrary to the situation with normal form games, the game does not consist of a single play, but jumps according to the probability measure $\pi(\cdot | \omega, s_\omega)$ to the next state and continues dynamically. In choosing an action in a certain state, a player not only takes into account the immediate payoff, but also his opportunities in future states.

Like in normal form games, the players are allowed to randomize their pure actions. A mixed strategy of player i in state ω is a probability distribution on S_ω^i . We identify the set of all probability distributions on S_ω^i with $\Sigma_\omega^i = \Delta(S_\omega^i)$. For $\sigma_\omega^i \in \Sigma_\omega^i$, the probability assigned to pure strategy $s_{\omega j}^i$ is given by $\sigma_{\omega j}^i$. The strategy space of the normal form game in state ω is therefore equal to $\Sigma_\omega = \times_{i \in N} \Sigma_\omega^i$. Given a mixed strategy combination $\sigma_\omega \in \Sigma_\omega$ and a strategy $\bar{\sigma}_\omega^i \in \Sigma_\omega^i$, we denote by $(\sigma_\omega^{-i}, \bar{\sigma}_\omega^i)$ the mixed strategy that results from replacing σ_ω^i by $\bar{\sigma}_\omega^i$. If a mixed strategy combination $\sigma_\omega \in \Sigma_\omega$ is played, then the instantaneous expected payoff of player i is denoted by $u^i(\omega, \sigma_\omega)$ and the expected transition to state $\bar{\omega}$ by $\pi(\bar{\omega} | \omega, \sigma_\omega)$.

At the different stages this randomization may depend on the history of the game up to that stage.

Definition 2.2. A *history* up to a stage k is a sequence $h^k = (\omega^0, s_{\omega^0}^0, \omega^1, s_{\omega^1}^0, \dots, \omega^{k-1}, s_{\omega^{k-1}}^0)$ that could have occurred up to time k , $k \geq 1$. Here ω^κ represents the state and s_{ω^κ} the actions of the players at stage κ , $\kappa = 0, \dots, k-1$.

Obviously, the set of histories up to time k equals $H^k = \times_{\kappa=0}^{k-1} H$, i.e., the k -fold Cartesian product of H .

Definition 2.3. A *behavior strategy* σ^i for player i is a sequence $\sigma^{i0}, \sigma^{i1}, \dots$ where $\sigma^{i0} \in \Sigma^i = \times_{\omega \in \Omega} \Sigma_\omega^i$ and $\sigma^{ik} : H^k \rightarrow \Sigma^i$ for all $k \geq 1$.

So, a *behavior strategy* σ^{ik} of player i specifies for each stage k , each state ω^k at time k , and each history h^k a probability distribution $\sigma^{ik}(h^k, \omega^k)$ on the action space $S_{\omega^k}^i$ of player i in state ω^k . Then $\sigma^{ik}(s_{\omega^k}^i | h^k, \omega^k)$ is the probability with which player i chooses action $s_{\omega^k}^i \in S_{\omega^k}^i$ at time k if state ω^k and history h^k have occurred.

Given initial state ω and strategy σ , the stream of expected payoffs is evaluated by

$$U^i(\omega, \sigma) := \sum_{k=0}^{\infty} \delta^k \cdot U^{ik}(\omega, \sigma),$$

where $U^{ik}(\omega, \sigma)$ denotes the expected instantaneous utility at stage k . Here, $U^i(\omega, \sigma)$ equals the total discounted expected payoff of player i when the discount factor equals δ , the starting state is ω and the strategy-tuple σ is played. Since the state and action spaces are assumed to be finite, $U^i(\omega, \sigma)$ exists.

Definition 2.4. A strategy-tuple σ is a *Nash equilibrium* if and only if σ^i is a best response to σ^{-i} for all $i \in N$.

If $n = 2$ and the payoffs are zero-sum, then the definition characterizes a minimax solution (see [32]).

3. Stationary equilibria

In the sequel of this paper we will restrict ourselves to stationary equilibria.¹ Before coming to the definition of a stationary equilibrium, first the concept of a stationary strategy has to be defined.

Definition 3.1. A *stationary strategy* ρ^i for player i is an element of Σ^i .

So, a *stationary strategy* for player i is a behavior strategy for which $\sigma^{ik}(h^k, \omega^k)$ is of the form $\sigma^i(\omega^k)$, i.e., a stationary strategy is a history and time independent strategy which depends on the state only. In the sequel, a stationary strategy for player i will be denoted by the symbol ρ^i . We define $\rho^i = (\rho_{\omega_1}^i, \dots, \rho_{\omega_z}^i)$, where ρ_{ω}^i is a probability measure on the action space S_{ω}^i for each $\omega \in \Omega$. So $\rho_{\omega}^i \in \Sigma_{\omega}^i$. If player i decides to play the stationary strategy $\rho^i \in \Sigma^i$, then every time that the system is in state ω , player i selects his pure action according to ρ_{ω}^i .

Suppose that ω is the initial state and the players decide to play a stationary strategy-tuple ρ . The total discounted expected payoff of player i is denoted by $U^i(\omega, \rho)$. The instantaneous payoff player i obtains in stage $k = 0$ equals $u^i(\omega, \rho_{\omega})$. The probability that at the next stage the state will be $\bar{\omega}$ equals $\pi(\bar{\omega} | \omega, \rho_{\omega})$. This gives rise to the following theorem (see e.g. [6]).

Theorem 3.2. When ρ is a stationary strategy-tuple and ω is the initial state, the expected payoffs are given by the following recursive formula:

$$U^i(\omega, \rho) = u^i(\omega, \rho) + \delta \cdot \sum_{\bar{\omega} \in \Omega} \pi(\bar{\omega} | \omega, \rho_{\omega}) U^i(\bar{\omega}, \rho).$$

Definition 3.3. A stationary strategy-combination $\rho \in \Sigma$ is a *stationary equilibrium* if it is a Nash equilibrium in stationary strategies.

It is a very well-known result that there exists a stationary best response to stationary strategies ρ^{-i} . Therefore, the following theorem holds.

Theorem 3.4. A stationary strategy-combination $\bar{\rho} \in \Sigma$ is a *stationary equilibrium* if for all i and all ω

$$U^i(\omega, \bar{\rho}) \geq U^i(\omega, \bar{\rho}^{-i}, \rho^i), \quad \rho^i \in \Sigma^i.$$

¹In fact, we restrict ourselves to a class which is even tighter. Namely, the one of perfect stationary equilibria which is a stationary equilibrium for all possible initial states.

Another famous result, see Fink [6,35] or [33], is the existence of a stationary equilibrium.

Theorem 3.5. *Every finite discounted stochastic game has a stationary equilibrium.*

The remainder of this section is devoted to the characterization of the set of stationary equilibria that is useful for numerical computations. Given that the other players play ρ^{-i} and the initial state is $\bar{\omega}$, player i faces the Markov decision problem of maximizing

$$U^i(\bar{\omega}, \rho^{-i}, \rho^i)$$

over ρ^i ,² subject to

$$\begin{aligned} \rho_{\omega j}^i &\geq 0, \quad (s_{\omega j}^i \in S_{\omega}^i, \omega \in \Omega), \\ \sum_{s_{\omega j}^i \in S_{\omega}^i} \rho_{\omega j}^i - 1 &= 0, \quad (\omega \in \Omega) \end{aligned}$$

for all possible initial states $\bar{\omega} \in \Omega$. If we define by $\mu_{\omega}^i(k)$ the present value of the total expected payoff for a system in state ω with k transitions remaining, we obtain the basic recurrence relation

$$\mu_{\omega}^i(k+1) = u^i(\omega, \rho_{\omega}) + \delta \cdot \sum_{\bar{\omega} \in \Omega} \pi(\bar{\omega} | \omega, \rho_{\omega}) \mu_{\bar{\omega}}^i(k)$$

or in vector form

$$\mu^i(k+1) = u^i(\rho) + \delta \Pi(\rho) \mu^i(k).$$

Note that³

$$\begin{aligned} \mu^i &= \lim_{k \rightarrow \infty} \mu^i(k) = \lim_{k \rightarrow \infty} \left(\sum_{\ell=0}^{k-1} (\delta \Pi(\rho))^{\ell} u^i(\rho) + (\delta \Pi(\rho))^k \mu^i(0) \right) \\ &= \sum_{\ell=0}^{\infty} (\delta \Pi(\rho))^{\ell} u^i(\rho) = [I - \delta \Pi(\rho)]^{-1} u^i(\rho). \end{aligned}$$

The vector μ^i may be called the vector of present values of player i , because each of its elements μ_{ω}^i is the present value of an infinite number of future expected payoffs discounted by the discount factor δ with ω the initial state.

²Here, the player maximizes over his stationary strategies rather than over his behavioral strategies. But as already mentioned there always exists a best response in stationary strategies when the opponents are playing stationary strategies. Therefore, the maximizing stationary strategies is also a maximizing behavioral strategy. So, it is not the case that nonstationary strategies are not allowed, they are not subject of selection.

³By Hadamard's theorem all eigenvalues of the matrix between the squared brackets have absolute value larger than 0. Thus zero is not an eigenvalue of that matrix and the inverse exists. Furthermore, the spectral radius of $\delta \Pi(\rho)$ is less than one. Therefore it holds that $[I - \delta \Pi(\rho)]^{-1} = \sum_{k=0}^{\infty} \delta^k \Pi(\rho)^k$. From this it is easily seen that the value of element (i, j) of the matrix $[I - \delta \Pi(\rho)]^{-1}$ gives the discounted expected number of times that the state is j when i is the initial state and ρ is played.

Because we are interested in the sequential decision process for large k (in fact for k equal to infinity), we substitute the present values $\mu_\omega^i = \lim_{k \rightarrow \infty} \mu_\omega^i(k)$ for the quantities $\mu_\omega^i(k)$ to obtain the recursive relation

$$\mu_\omega^i = u^i(\omega, \rho_\omega) + \delta \cdot \sum_{\bar{\omega} \in \Omega} \pi(\bar{\omega} | \omega, \rho_\omega) \mu_{\bar{\omega}}^i.$$

Given that the other players play ρ^{-i} , player i maximizes μ_ω^i subject to $\rho_\omega^i \in \Sigma_\omega^i$ for all $\omega \in \Omega$.

Because player i maximizes each μ_ω^i for $\omega \in \Omega$, he also maximizes $\sum_{\omega \in \Omega} \mu_\omega^i$. Given that the other players play ρ^{-i} , player i maximizes

$$\sum_{\omega \in \Omega} \mu_\omega^i$$

subject to

$$\begin{aligned} u^i(\omega, \rho_\omega^{-i}, \rho_\omega^i) + \delta \cdot \sum_{\bar{\omega} \in \Omega} \pi(\bar{\omega} | \omega, \rho_\omega^{-i}, \rho_\omega^i) \mu_{\bar{\omega}}^i - \mu_\omega^i &= 0 \quad (\omega \in \Omega), \\ \rho_{\omega j}^i &\geq 0 \quad (s_{\omega j}^i \in S_\omega^i, \omega \in \Omega), \\ \sum_{s_{\omega j}^i \in S_\omega^i} \rho_{\omega j}^i - 1 &= 0 \quad (\omega \in \Omega). \end{aligned}$$

The necessary and sufficient first-order conditions for this maximization problem are

$$\begin{aligned} \alpha_\omega^i \{ u^i(\omega, \rho_\omega^{-i}, s_{\omega j}^i) + \delta \cdot \sum_{\bar{\omega} \in \Omega} \pi(\bar{\omega} | \omega, \rho_\omega^{-i}, s_{\omega j}^i) \mu_{\bar{\omega}}^i \} + \beta_{\omega j}^i - \gamma_\omega^i &= 0 \\ (s_{\omega j}^i \in S_\omega^i, \omega \in \Omega), \\ \beta_{\omega j}^i &\geq 0, \quad \rho_{\omega j}^i \geq 0, \quad \beta_{\omega j}^i \rho_{\omega j}^i = 0 \quad (s_{\omega j}^i \in S_\omega^i, \omega \in \Omega), \\ \sum_{s_{\omega j}^i \in S_\omega^i} \rho_{\omega j}^i - 1 &= 0 \quad (\omega \in \Omega), \\ u^i(\omega, \rho_\omega^{-i}, \rho_\omega^i) + \delta \cdot \sum_{\bar{\omega} \in \Omega} \pi(\bar{\omega} | \omega, \rho_\omega^{-i}, \rho_\omega^i) \mu_{\bar{\omega}}^i - \mu_\omega^i &= 0 \quad (\omega \in \Omega), \\ \alpha^i [I - \delta \Pi(\rho^{-i}, \rho^i)] &= \mathbb{1}. \end{aligned}$$

Here, α , β and γ are the Lagrange multipliers of the first, second and third set of constraints, and $\mathbb{1}$ denotes the vector containing ones only. Then, by multiplying by $\rho_{\omega j}^i$ and summing over j , for ρ^i a best response to ρ^{-i} ,

$$\begin{aligned} 0 &= \alpha_\omega^i \{ u^i(\omega, \rho_\omega^{-i}, \rho_\omega^i) + \delta \cdot \sum_{\bar{\omega} \in \Omega} \pi(\bar{\omega} | \omega, \rho_\omega^{-i}, \rho_\omega^i) \mu_{\bar{\omega}}^i \} \\ &\quad + \sum_{s_{\omega j}^i \in S_\omega^i} \beta_{\omega j}^i \rho_{\omega j}^i - \sum_{s_{\omega j}^i \in S_\omega^i} \gamma_\omega^i \rho_{\omega j}^i \\ &= \alpha_\omega^i \mu_\omega^i - \gamma_\omega^i \end{aligned}$$

for all $\omega \in \Omega$. Therefore, since $\alpha^i = [I - \delta \Pi(\rho^{-i}, \rho^i)]^{-1} \mathbb{1} = \frac{1}{1-\delta} \mathbb{1} \gg 0$,⁴

$$\frac{\gamma_\omega^i}{\alpha_\omega^i} = \mu_\omega^i.$$

After division by α_ω^i , the following necessary and sufficient conditions remain, where $\lambda_{\omega j}^i$ is defined as the ratio of $\beta_{\omega j}^i$ and α_ω^i :

$$\begin{aligned} u^i(\omega, \rho_\omega^{-i}, s_{\omega j}^i) + \delta \cdot \sum_{\bar{\omega} \in \Omega} \pi(\bar{\omega} | \omega, \rho_\omega^{-i}, s_{\omega j}^i) \mu_{\bar{\omega}}^i + \lambda_{\omega j}^i - \mu_\omega^i &= 0 \\ (s_{\omega j}^i \in S_\omega^i, \omega \in \Omega), \\ \lambda_{\omega j}^i \geq 0, \quad \rho_{\omega j}^i \geq 0, \quad \lambda_{\omega j}^i \rho_{\omega j}^i &= 0 \quad (s_{\omega j}^i \in S_\omega^i, \omega \in \Omega), \\ \sum_{s_{\omega j}^i \in S_\omega^i} \rho_{\omega j}^i - 1 &= 0 \quad (\omega \in \Omega). \end{aligned}$$

Here, $\lambda_{\omega j}^i$ is the shadowprice of playing strategy $s_{\omega j}^i$, i.e., the disutility from a one-shot deviation at $t = 0$ of a marginal increase in the probability $\rho_{\omega j}^i$ by which pure strategy $s_{\omega j}^i$ is played, and μ_ω^i is the expected payoff of player i when the initial state is ω , ρ^{-i} is played by his opponents, and player i chooses a best response. The last equality, $\sum_{s_{\omega j}^i \in S_\omega^i} \rho_{\omega j}^i - 1 = 0$, makes sure that ρ_ω^i is a member of Σ_ω^i for all $\omega \in \Omega$.

Since for a stationary equilibrium it holds that a strategy-tuple constitutes mutually best responses, we have found that the set of stationary equilibria can be fully characterized by the system of equalities and inequalities in Theorem 3.6.

Theorem 3.6. *A strategy $\rho \in \Sigma$ is a stationary equilibrium if and only if it is part of a solution to*

$$\begin{aligned} u^i(\omega, \rho_\omega^{-i}, s_{\omega j}^i) + \delta \cdot \sum_{\bar{\omega} \in \Omega} \pi(\bar{\omega} | \omega, \rho_\omega^{-i}, s_{\omega j}^i) \mu_{\bar{\omega}}^i + \lambda_{\omega j}^i - \mu_\omega^i &= 0 \\ (s_{\omega j}^i \in S_\omega^i, \omega \in \Omega, i \in N), \\ \lambda_{\omega j}^i \geq 0, \quad \rho_{\omega j}^i \geq 0, \quad \lambda_{\omega j}^i \rho_{\omega j}^i &= 0 \quad (s_{\omega j}^i \in S_\omega^i, \omega \in \Omega, i \in N), \\ \sum_{s_{\omega j}^i \in S_\omega^i} \rho_{\omega j}^i - 1 &= 0 \quad (\omega \in \Omega, i \in N). \end{aligned}$$

The system above suggests that only one-shot deviations have to be considered. We explain intuitively the reason of this surprising phenomenon often referred to as the one-deviation property.

Suppose the other players play stationary strategies ρ^{-i} and suppose that ρ^i is the best stationary response of player i . If player i is not able to improve his utility by a

⁴Since the value of element (i, j) of the matrix $[I - \delta \Pi(\rho)]^{-1}$ gives the discounted expected number of times that the state is j when i is the initial state and ρ is played, the rows sum up to $\frac{1}{1-\delta}$, which is larger than 0 (even large than 1).

Fig. 1. Example to show that the expected utility property does not hold.

deviation to his strategy ρ^i in one stage only, then it follows by a backward induction argument, that neither finitely many deviations to his strategy will make player i better off. Suppose now that player i can improve his utility by infinitely many changes. Then, by a profit-to-go argument, player i can also increase his payoff by finitely many changes, which is not possible.⁵

We end this section with an example in which it is shown that expected utility in stationary strategies does not hold for the class of stochastic games. This causes a number of technical difficulties for the convergence proof of our algorithm.

Example 3.7. The stochastic game of Fig. 1 is a game in which there are two states p_H and p_L and one player. In each state the player can choose between actions q_H and q_L . In the upper-left corner of each square the intermediate payoff of the player is displayed. In the lower-right corner the transition probabilities are given. So, if the player chooses q_H when the state is p_H , then he earns 1 today and the next state will be p_L with probability 1. If the player chooses q_L when the state is p_L , then he earns 1 today and the next state will be p_H for sure. Future payoffs are discounted by a factor $1/2$.

Consider the strategies $\rho^- = (q_L, q_H)$ and $\rho^+ = (q_H, q_L)$, where the first argument is the strategy when the state is p_H and the second when the state is p_L . Then $U(p_L, \rho^-) = U(p_H, \rho^-) = 1$, $U(p_L, \rho^+) = 2/3$ and $U(p_H, \rho^+) = 4/3$.

However,

$$U(p_H, \frac{1}{2}\rho^- + \frac{1}{2}\rho^+) = \frac{6}{5} \neq \frac{7}{6} = \frac{1}{2} \cdot 1 + \frac{1}{2} \cdot \frac{4}{3} = \frac{1}{2} U(p_H, \rho^-) + \frac{1}{2} U(p_H, \rho^+),$$

$$U(p_L, \frac{1}{2}\rho^- + \frac{1}{2}\rho^+) = \frac{4}{5} \neq \frac{5}{6} = \frac{1}{2} \cdot 1 + \frac{1}{2} \cdot \frac{2}{3} = \frac{1}{2} U(p_L, \rho^-) + \frac{1}{2} U(p_L, \rho^+).$$

So, expected utility does not hold.

Next we provide an example that shows that even for a stochastic game with two players the ordered field property does not hold, i.e. even if all exogenous variables

⁵Define $M = \max_{(i, \omega, s_{\omega})} |u^i(\omega, s_{\omega})|$. Then the maximum payoff a player can earn from time k on is bounded from above by $\delta^k (1 + \delta + \delta^2 + \dots) M = \frac{\delta^k}{1-\delta} M$, the so-called maximum ‘profit-to-go’ value. Suppose player i is able to improve his utility by ε by means of infinitely many changes. When k grows large the profit-to-go value is at a certain point less than ε (this is when $k > \log(\frac{\varepsilon(1-\delta)}{M}) / \log(\delta)$). This means that the utility improvement by changes until time k (finitely many changes) was positive.

of the game are rational numbers, it is still possible that equilibrium strategies involve irrational numbers. In the literature on the computation of Nash equilibria in normal form games, a distinction is made between 2-player games on the one hand and 3 or more players on the other hand. For the class of 2-player games exact algorithms are possible, because of the bilinear structure of such games (see, for instance the algorithm of Lemke and Howson [18]). For stochastic games this distinction disappears. The system of equations of Theorem 3.6 is not bilinear, even for 2-player games.

Example 3.8. The stochastic game of Fig. 2 originates from Parthasarathy and Raghavan [28]. The game has two states. In the first state player 1 is the row player and has two pure actions and player 2 is the column player and also has two pure actions. In the second state both players have one action only. The transition probabilities (displayed in the right-lower corners) show that once state 2 is reached, it will never leave this state. The numbers in the left-upper corners specify the amount player 2 has to pay to player 1 (this game is of the zero-sum type).

The only equilibrium in this stochastic game is both players playing (in state 1 of course) the first action with probability

$$\frac{3 + \frac{1}{3}(\sqrt{13} - 2)}{4 + \frac{2}{3}(\sqrt{13} - 2)}$$

and the second with probability

$$\frac{1 + \frac{1}{3}(\sqrt{13} - 2)}{4 + \frac{2}{3}(\sqrt{13} - 2)}.$$

So, the data of the game and the solution to the game are not in the same Archimedean field.

Although that this 2-player stochastic game has only rational numbers of payoffs and transition probabilities, the unique Nash equilibrium involves strategies with irrational probabilities. This means that the ordered field property does not hold. A straightforward application of a Lemke–Howson type algorithm to 2-player stochastic games is therefore not possible. This shows one more time that the class of stochastic games is considerably more difficult than the class of normal form games.

Fig. 2. Example to show that the ordered field property does not hold.

4. The stochastic tracing procedure

The linear tracing procedure as presented in Harsanyi and Selten [10] models a process of convergent expectations by which rational players will come to adopt, and expect each other to adopt, a particular equilibrium as a solution for a given game. Before applying the tracing procedure, every player is assumed to have a subjective probability distribution expressing his expectation about the strategic choices of the other players. Each player is assumed to use the same theory to determine his subjective probability distributions, which makes that all players have the same expectations about the other players. This common subjective probability distribution is called the prior. In the naive Bayesian approach, all players choose best responses to their prior beliefs and would in this way reach a strategy-combination that does not constitute an equilibrium in general. In the linear tracing procedure, the information on the best responses is only gradually fed back into the expectations of the players. As the linear tracing procedure proceeds, both the priors and their best responses will gradually change until both converge to some equilibrium of the game.

In Harsanyi and Selten [10] the linear tracing procedure is defined for normal form games and for extensive form games with a perfect recall information structure. For a normal form game $\Gamma = \langle N, \{S^i\}_{i \in N}, \{u^i\}_{i \in N} \rangle$ and a prior $p \in \Sigma$ the linear tracing procedure is defined by tracing a curve in the set of Nash equilibria of the games $\Gamma^t = \langle N, \{S^i\}_{i \in N}, \{\tilde{u}^i(t)\}_{i \in N} \rangle$ for $t \in [0, 1]$, where $\tilde{u}^i(t, s) = (1 - t)u^i(p^{-i}, s^i) + tu^i(s)$. For extensive form games, they first transform the game into one in standard form and subsequently define the tracing procedure for that class of games. They did not define the tracing procedure for stochastic games, which are games with instantaneous payoffs and infinite time horizon. Since expected utility in stationary strategies does not hold in stochastic games (see Example 3.7) it is not possible to transform this game into one in standard form. The extension of the tracing procedure to the class of stochastic games is far from obvious.

There are at least four ways to extend the tracing procedure of Harsanyi and Selten to the setting of stochastic games. Choices have to be made whether a player holds correlated beliefs within a state or not, and whether a player holds correlated beliefs across time or not. For the extension of the linear tracing procedure to stochastic games that we study in this paper, we assume that beliefs are correlated within states and that they are not correlated across time.

Correlation within states means that when a player knows that some opponent plays according to the prior (which he expects with probability $1 - t$), he expects all opponents to play according to the prior. This is equivalent to the way Harsanyi and Selten define the tracing procedure for normal form games.

Absence of correlation across time means that even when a player knows that his opponents are playing according to the prior today, these opponents might not play according to the prior in future stages. In all future events he faces independent lotteries which assigns probability $1 - t$ to play against the prior strategies of his opponents. Assuming that beliefs are not correlated across time captures the

assumption of stationarity. The beliefs of a player depend only on the state reached and not on the time at which it is reached. A more extensive discussion on the definition of the stochastic tracing procedure is the issue of Herings and Peeters [15].

Fix some stochastic game Γ and some prior $p \in \Sigma$. For every $t \in [0, 1]$, the stochastic tracing procedure generates a stationary equilibrium of the stochastic game

$$\Gamma^t = \langle N, \Omega, \{S_\omega^i\}_{\omega \in \Omega, i \in N}, \{\tilde{u}^i(t)\}_{i \in N}, \{\tilde{\pi}^i(t)\}_{i \in N}, \delta \rangle,$$

where the instantaneous payoff function $\tilde{u}^i(t) : H \rightarrow \mathbb{R}$ of player i is defined by

$$\tilde{u}^i(t; \omega, s_\omega) = (1 - t)u^i(\omega, p_\omega^{-i}, s_\omega^i) + t\bar{u}^i(\omega, s_\omega)$$

and where the transition mapping $\tilde{\pi}^i(t)$ is defined by

$$\tilde{\pi}^i(t; \bar{\omega} \mid \omega, s_\omega) = (1 - t)\pi(\bar{\omega} \mid \omega, p_\omega^{-i}, s_\omega^i) + t\pi(\bar{\omega} \mid \omega, s_\omega).$$

Note that $\tilde{\pi}^i(t)$ may be different for different players when $t < 1$. This is consistent with the fact that the tracing procedure should be thought of as a reasoning process. The mapping $\tilde{\pi}^i(t)$ should be thought of as what the players think that the transition probabilities are in the stochastic game Γ^t . The expected payoff of player i is easily shown to satisfy the recursive relation

$$\tilde{U}^i(t; \omega, \rho) = \tilde{u}^i(t; \omega, \rho_\omega) + \delta \cdot \sum_{\bar{\omega} \in \Omega} \tilde{\pi}^i(t; \bar{\omega} \mid \omega, \rho_\omega) \tilde{U}^i(t; \bar{\omega}, \rho).$$

The stochastic game Γ^0 corresponds to a trivial stochastic game, where all players believe that all their opponents play with probability 1 according to the prior belief. The stochastic game Γ^1 coincides with the original stochastic game Γ . A best response against a strategy combination $\rho^{-i} \in \Sigma^{-i}$ in the stochastic game Γ^t corresponds to a best response against the stationary probability distribution $(1 - t)[p^{-i}] + t[\rho^{-i}]$ on S^{-i} in the stochastic game Γ .

The stochastic tracing procedure $\mathcal{S}(\Gamma, p)$ is defined as the set of pairs (t, ρ) for which it holds that ρ is a stationary equilibrium of the stochastic game Γ^t , i.e.,

$$\mathcal{S}(\Gamma, p) = \{(t, \rho) \in [0, 1] \times \Sigma \mid \rho \text{ is a stationary equilibrium of } \Gamma^t\}.$$

The stochastic tracing procedure is said to be feasible if there exists a path in $\mathcal{S}(\Gamma, p)$ connecting a best response against the prior to a stationary equilibrium of the stochastic game Γ , i.e., there exists a continuous function $\gamma : [0, 1] \rightarrow \mathcal{S}(\Gamma, p)$ such that $\gamma(0) \in \mathcal{S}(\Gamma, p) \cap (\{0\} \times \Sigma)$ and $\gamma(1) \in \mathcal{S}(\Gamma, p) \cap (\{1\} \times \Sigma)$. In general there may be many trajectories $\gamma([0, 1])$ that link a stationary equilibrium of Γ^0 to a stationary equilibrium of Γ^1 . If this trajectory is unique, then the stochastic tracing procedure is said to be well-defined. If the stochastic tracing procedure is well-defined, then it selects a unique stationary equilibrium of the stochastic game Γ .

For a simple proof of the feasibility of the linear tracing procedure for normal form games see Herings [12] and for the well-definedness see Herings and Peeters [14]. The derivation of such properties for the stochastic tracing procedure is part of present paper.

5. Structure of the stochastic tracing procedure

The size of any stochastic game Γ can be characterized by a vector $\eta = (n, z, \{m_\omega^i\}_{\omega=1, \dots, z}^{i=1, \dots, n})$, specifying the number of players, the number of states, and the number of pure strategies available to a player in a state. We will call η the size vector of a stochastic game Γ . For any possible size vector η , transition map π consistent with η and discount factor δ , the set $\mathcal{G}(\eta, \pi, \delta)$ of all stochastic games Γ possessing η as their size vector, having π as their transition map, and δ as discount factor, is called the size class generated by (η, π, δ) .

Every stochastic game Γ in the size class $\mathcal{G}(\eta, \pi, \delta)$ is characterized by a vector $u(\Gamma)$ that contains exactly nm payoffs, where $m = |S| = \prod_{(i, \omega) \in N \times \Omega} m_\omega^i$. We identify a stochastic game Γ with the vector $u(\Gamma)$, and we identify the size class $\mathcal{G}(\eta, \pi, \delta)$ with the set of all possible real vectors of size nm , that is, with an nm -dimensional Euclidean space. It is now possible to define the distance $\varrho(\Gamma, \Gamma')$ between two stochastic games and the Lebesgue measure of a set of stochastic games.

A given mathematical statement $\mathcal{Z}$ is said to be true for an open set of (almost all, set of generic) stochastic games if, for every possible size class $\mathcal{G}(\eta, \pi, \delta)$, the set $\tilde{\mathcal{G}}(\mathcal{Z})$ of all stochastic games Γ in $\mathcal{G}(\eta, \pi, \delta)$ for which the statement $\mathcal{Z}$ is true is open (has full Lebesgue measure, is open and of full Lebesgue measure) with respect to $\mathcal{G}(\eta, \pi, \delta)$.

Fix η , π and δ .⁶ Let a subset B^* of $S^* = \bigcup_{(i, \omega) \in N \times \Omega} S_\omega^i$ be given with the property that for every player i and for every state ω there is at least one pure strategy $s_{\omega j}^i$ in B^* , so $B_\omega^i = B^* \cap S_\omega^i \neq \emptyset$, for every player i , for every state ω . Such a set B^* is called admissible. The sets B^* are used to decompose $\mathcal{S}(\Gamma, p)$ in subsets $\mathcal{S}(\Gamma, p, B^*)$, each having a differentiable manifold structure. The set $\mathcal{S}(\Gamma, p, B^*)$ contains those elements of $\mathcal{S}(\Gamma, p)$ where only strategies in B^* are played with positive probability. It is defined by

$$\mathcal{S}(\Gamma, p, B^*) = \left\{ (t, \rho) \in \mathcal{S}(\Gamma, p) \left| \begin{array}{l} s_{\omega j}^i \notin B^* \Rightarrow \rho_{\omega j}^i = 0 \\ s_{\omega j}^i \in B^* \Rightarrow s_{\omega j}^i \in \underset{s_{\omega \ell}^i \in S_\omega^i}{\operatorname{argmax}} \tilde{U}^i(t; \omega, \rho^{-i}, \rho_{-\omega}^i, s_{\omega \ell}^i) \end{array} \right. \right\}.$$

It follows that

$$\mathcal{S}(\Gamma, p) = \bigcup_{B^*} \mathcal{S}(\Gamma, p, B^*).$$

Two sets $\mathcal{S}(\Gamma, p, B^*)$ and $\mathcal{S}(\Gamma, p, \tilde{B}^*)$ can only have a point (t, ρ) in common if there is a player i and a strategy $s_{\omega j}^i$ such that $\rho_{\omega j}^i = 0$ and $s_{\omega j}^i \in \underset{s_{\omega \ell}^i \in S_\omega^i}{\operatorname{argmax}} \tilde{U}^i(t; \omega, \rho^{-i}, \rho_{-\omega}^i, s_{\omega \ell}^i)$, so $s_{\omega j}^i$ is a best response to $(\rho^{-i}, \rho_{-\omega}^i)$ that is played with probability zero.

⁶It is possible to fix only η , and to define open sets of (almost all, generic sets of) stochastic games if these properties hold for every possible size class $\mathcal{G}(\eta)$. All forthcoming proofs can be readily adapted to get results corresponding to these definitions.

To analyze the structure of $\mathcal{S}(\Gamma, p)$ and the sets $\mathcal{S}(\Gamma, p, B^*)$, we design systems of equalities and inequalities whose solutions characterize these sets. By Theorem 3.6, an element (t, ρ) belongs to $\mathcal{S}(\Gamma, p)$ if and only if it is part of a solution to

$$\begin{aligned} & \tilde{u}^i(t; \omega, \rho_\omega^{-i}, s_{\omega j}^i) + \delta \cdot \sum_{\bar{\omega} \in \Omega} \tilde{\pi}^i(t; \bar{\omega} \mid \omega, \rho_\omega^{-i}, s_{\omega j}^i) \mu_{\bar{\omega}}^i + \lambda_{\omega j}^i - \mu_\omega^i = 0 \\ & (s_{\omega j}^i \in S_\omega^i, \omega \in \Omega, i \in N), \\ & \lambda_{\omega j}^i \geq 0, \quad \rho_{\omega j}^i \geq 0, \quad \lambda_{\omega j}^i \rho_{\omega j}^i = 0 \quad (s_{\omega j}^i \in S_\omega^i, \omega \in \Omega, i \in N), \\ & \sum_{s_{\omega j}^i \in S_\omega^i} \rho_{\omega j}^i - 1 = 0 \quad (\omega \in \Omega, i \in N). \end{aligned}$$

Given an admissible subset B^* , we define $\mathcal{O}(\Gamma, p, B^*)$ as the set of solutions (t, ρ, λ, μ) to the following system of equalities and inequalities:

$$\begin{aligned} (1) \quad & \tilde{u}^i(t; \omega, \rho_\omega^{-i}, s_{\omega j}^i) + \delta \cdot \sum_{\bar{\omega} \in \Omega} \tilde{\pi}^i(t; \bar{\omega} \mid \omega, \rho_\omega^{-i}, s_{\omega j}^i) \mu_{\bar{\omega}}^i + \lambda_{\omega j}^i - \mu_\omega^i = 0, \\ & (s_{\omega j}^i \in S_\omega^i, \omega \in \Omega, i \in N), \\ (2) \quad & \rho_{\omega j}^i = 0 \quad (s_{\omega j}^i \notin B_\omega^i, \omega \in \Omega, i \in N), \\ (3) \quad & \lambda_{\omega j}^i = 0 \quad (s_{\omega j}^i \in B_\omega^i, \omega \in \Omega, i \in N), \\ (4) \quad & \sum_{s_{\omega j}^i \in S_\omega^i} \rho_{\omega j}^i - 1 = 0 \quad (\omega \in \Omega, i \in N), \\ (5) \quad & \rho_{\omega j}^i \geq 0 \quad (s_{\omega j}^i \in B_\omega^i, \omega \in \Omega, i \in N), \\ (6) \quad & \lambda_{\omega j}^i \geq 0 \quad (s_{\omega j}^i \notin B_\omega^i, \omega \in \Omega, i \in N), \\ (7) \quad & t \geq 0, \\ (8) \quad & 1 - t \geq 0. \end{aligned}$$

Every solution of system (1)–(8) is a solution of the system in Theorem 3.6, hence an equilibrium.

The fact that for stochastic games the system of equalities and inequalities differs from the case of normal form games, is not the only difficulty. Since expected utility does not hold, we cannot use any property that is derived from it. In particular, it is not even obvious that there is a unique best response to the prior in pure stationary strategies.

Theorem 5.1 implies that $(t, \rho) \in \mathcal{S}(\Gamma, p, B^*)$ if and only if there exists $\lambda \in \mathbb{R}^{m^*}$ and $\mu \in \mathbb{R}^{nz}$ such that equalities (1)–(4) and inequalities (5)–(8) are satisfied, where $m^* = |S^*| = \sum_{(i, \omega) \in N \times \Omega} m_\omega^i$.

Theorem 5.1. *Let a stochastic game $\Gamma \in \mathcal{G}(\eta, \pi, \delta)$ and a prior $p \in \Sigma$ be given. For all admissible subsets B^* of S^* , the sets $\mathcal{S}(\Gamma, p, B^*)$ and $\mathcal{O}(\Gamma, p, B^*)$ are C^∞ diffeomorphic.*

Proof. Let B^* be an admissible subset of S^* . For every $i \in N$ and $\omega \in \Omega$ take an element $\bar{s}_\omega^i \in B^*$. Define a function $f: [0, 1] \times \Sigma \rightarrow \mathbb{R} \times \mathbb{R}^{m^*} \times \mathbb{R}^{m^*} \times \mathbb{R}^{nz}$ by $f(t, \rho) =$

(t, ρ, λ, μ) , where (in vector form with a slight abuse of notation) $\mu^i = [I - \delta \cdot \tilde{\Pi}^i(t; \rho^{-i}, \bar{s}^i)]^{-1} \bar{u}^i(t; \rho^{-i}, \bar{s}^i)$ and $\lambda_{\omega\ell}^i = \mu_{\omega}^i - \bar{u}^i(t; \omega, \rho_{\omega}^{-i}, s_{\omega\ell}^i) - \delta \cdot \sum_{\bar{\omega} \in \Omega} \bar{\pi}^i(t; \bar{\omega} | \omega, \rho_{\omega}^{-i}, s_{\omega\ell}^i) \mu_{\bar{\omega}}^i$. Then $f(t, \rho) \in \mathcal{O}(\Gamma, p, B^*)$ if and only if $(t, \rho) \in \mathcal{S}(\Gamma, p, B^*)$. Note that f defined in this way is a C^∞ diffeomorphism. $\square$

From Theorem 5.1 it follows that for all $(t, \rho) \in \mathcal{S}(\Gamma, p, B^*)$ there is a unique λ and a unique μ such that $(t, \rho, \lambda, \mu) \in \mathcal{O}(\Gamma, p, B^*)$. Vice versa, for all $(t, \rho, \lambda, \mu) \in \mathcal{O}(\Gamma, p, B^*)$ it holds that $(t, \rho) \in \mathcal{S}(\Gamma, p, B^*)$.

The analysis of the system of equalities and inequalities (1)–(8) provides the following result.

Theorem 5.2. *There exists an open set of stochastic games and priors $(\Gamma, p) \in \mathcal{G}(\eta, \pi, \delta) \times \Sigma$ with full Lebesgue measure such that for all admissible subsets B^* of S^* , $\mathcal{S}(\Gamma, p, B^*)$ is a compact 1-dimensional C^∞ manifold with boundary. Moreover, a boundary point (t, ρ) of $\mathcal{S}(\Gamma, p, B^*)$ is either*

- (i) *not a boundary point of $\mathcal{S}(\Gamma, p, \bar{B}^*)$ for all $\bar{B}^* \neq B^*$ and lies in $\{0, 1\} \times \Sigma$, or*
- (ii) *is a boundary point of exactly one $\mathcal{S}(\Gamma, p, \bar{B}^*)$ with $\bar{B}^* \neq B^*$ and belongs to $(0, 1) \times \Sigma$. Moreover, B^* and $\bar{B}^*$ differ in exactly one element, say $s_{\omega j}^i$, for which $\rho_{\omega j}^i = 0$ and $s_{\omega j}^i$ is a best response to ρ in Γ^i .*

Proof. The proof of this theorem is analogous to the proofs of Theorems 2 and 3 of Herings and Peeters [14], and involves three lemmas. The only thing left to do is to prove those lemmas for the stochastic game situation, which is done in Appendix A. The inference is then analogously to the proofs of Herings and Peeters [14], since it only uses the structural manifold properties of the lemmas.

In Lemma A.1, we prove that the set of solutions to (1)–(8) is 1-dimensional if the inequalities are strict. Then in Lemma A.2, we prove that the set of solutions to (1)–(8) is zero-dimensional if exactly one of the inequalities is binding. Finally, it is proved in Lemma A.3 that there is no solution to (1)–(8) with more than one binding inequality.

It follows that $\mathcal{O}(\Gamma, p, B^*)$ is a 1-dimensional manifold with boundary. A point in $\mathcal{O}(\Gamma, p, B^*)$ is a boundary point if and only if exactly one of the inequalities in (5)–(8) holds with equality. By Theorem 5.1 these properties carry over to $\mathcal{S}(\Gamma, p, B^*)$.

It is easily seen that for a boundary point alternative (i) holds when the binding inequality comes from (7) or (8) and that alternative (ii) holds when the binding inequality comes from (5) or (6). $\square$

This theorem implies that, for almost every Γ and p , for all admissible sets B^* of S^* the set $\mathcal{S}(\Gamma, p, B^*)$ consists of a finite number of smooth arcs and loops.⁷ Each arc has two boundary points. The structure of $\mathcal{S}(\Gamma, p, B^*)$ is therefore a simple one; all

⁷ Let X be a topological space. It is called an arc if it is homeomorphic to the closed unit interval $[0, 1]$ and it has two boundary points in this case; it is called a loop if it is homeomorphic to the unit circle in $\mathbb{R}^2$ and it has no boundary points in that case.

kinds of complications like bifurcations, spirals, higher dimensional solution sets, diverging behavior, etc. are excluded.

If $\mathcal{S}(\Gamma, p, B^*)$ has a boundary point in $(0, 1) \times \Sigma$, then there is a unique admissible subset $\bar{B}^* \neq B^*$ such that $\mathcal{S}(\Gamma, p, \bar{B}^*)$ has this boundary point in common with $\mathcal{S}(\Gamma, p, B^*)$. The cardinality of $\bar{B}^*$ is one less or one greater than the cardinality of B^* , depending on whether in the common boundary point a strategy in B^* is played with probability zero, or whether a strategy not in B^* is a best response. The uniqueness of the set $\bar{B}^*$ implies that the sets $\mathcal{S}(\Gamma, p, B^*)$ and $\mathcal{S}(\Gamma, p, \bar{B}^*)$ are nicely linked to each other.

If $\mathcal{S}(\Gamma, p, B^*)$ has a boundary point in $\{0, 1\} \times \Sigma$, then this point does not belong to any other set $\mathcal{S}(\Gamma, p, \bar{B}^*)$. This implies that such a boundary point is also a boundary point of $\mathcal{S}(\Gamma, p)$.

Formally, the structure of $\mathcal{S}(\Gamma, p)$ is as follows.

Theorem 5.3. *There exists an open set of stochastic games and priors $(\Gamma, p) \in \mathcal{G}(\eta, \pi, \delta) \times \Sigma$ with full Lebesgue measure such that $\mathcal{S}(\Gamma, p)$ is a compact 1-dimensional piecewise C^∞ manifold with boundary. The boundary of $\mathcal{S}(\Gamma, p)$ is given by the intersection of $\mathcal{S}(\Gamma, p)$ and $\{0, 1\} \times \Sigma$. There is a unique boundary point in $\{0\} \times \Sigma$ which corresponds to a strategy-combination in pure strategies.*

Proof. The intuition of the proof of the first part of this theorem is given in the text above the theorem. A rigorous proof can be given using the Lemke-Howson argument. For a rigorous development of this argument, see for instance Herings and Peeters [14].

The second part of the theorem, that there is generically a unique boundary point in $\{0\} \times \Sigma$ and that this boundary point is in pure strategies, remains to be shown.

Suppose there is a best response $\bar{\rho}^i$ to the prior in mixed strategies. Then for some state ω player i plays under $\bar{\rho}^i$ at least two pure strategies with positive probability mass, say $s_{\omega 1}^1$ and $s_{\omega 2}^1$. Since $\bar{\rho}^i$ is an optimal strategy it is part of a solution to the necessary and sufficient conditions above Theorem 3.6. So,

$$\lambda_{\omega 1}^1 = \lambda_{\omega 2}^1 = 0$$

and

$$\begin{aligned} u^i(\omega, p_\omega^{-i}, s_{\omega 1}^i) + \delta \cdot \sum_{\bar{\omega} \in \Omega} \pi(\bar{\omega} \mid \omega, p_\omega^{-i}, s_{\omega 1}^i) \mu_{\bar{\omega}}^i &= \mu_\omega^i \\ &= u^i(\omega, p_\omega^{-i}, s_{\omega 2}^1) + \delta \cdot \sum_{\bar{\omega} \in \Omega} \pi(\bar{\omega} \mid \omega, p_\omega^{-i}, s_{\omega 2}^1) \mu_{\bar{\omega}}^i. \end{aligned}$$

It follows that every strategy ρ^i with $\rho_{\omega 1}^i, \rho_{\omega 2}^i \geq 0$ and $\rho_{-\{\omega 1, \omega 2\}}^i = \bar{\rho}_{-\{\omega 1, \omega 2\}}^i$ is a best response to the prior. This contradicts the local finiteness of the number of solutions at $t = 0$. Therefore, generically, all best responses are in pure strategies.

Suppose there are two pure best responses $\bar{\rho}^i$ and $\hat{\rho}^i$ against the prior. Then $\bar{\mu}_\omega^i = \hat{\mu}_\omega^i$ for all $\omega \in \Omega$. Take $\rho^i = \frac{1}{2} \bar{\rho}^i + \frac{1}{2} \hat{\rho}^i$, $\mu^i = \bar{\mu}^i$, and $\lambda_{\omega j}^i = 0$ if $\rho_{\omega j}^i > 0$ and $\lambda_{\omega j}^i = \bar{\lambda}_{\omega j}^i$ otherwise. Now indeed the complementarity conditions are satisfied and ρ^i is a member of Σ^i .

Suppose $\rho_{\omega j}^i > 0$. Then it holds that $\bar{\rho}_{\omega j}^i > 0$ or $\hat{\rho}_{\omega j}^i > 0$; suppose without loss of generality that $\bar{\rho}_{\omega j}^i > 0$. Then

$$\begin{aligned} 0 &= u^i(\omega, p_{\omega}^{-i}, s_{\omega j}^i) + \delta \cdot \sum_{\bar{\omega} \in \Omega} \pi(\bar{\omega} \mid \omega, p_{\omega}^{-i}, s_{\omega j}^i) \bar{\mu}_{\bar{\omega}}^i - \bar{\mu}_{\omega}^i \\ &= u^i(\omega, p_{\omega}^{-i}, s_{\omega j}^i) + \delta \cdot \sum_{\bar{\omega} \in \Omega} \pi(\bar{\omega} \mid \omega, p_{\omega}^{-i}, s_{\omega j}^i) \mu_{\bar{\omega}}^i + \lambda_{\omega j}^i - \mu_{\omega}^i. \end{aligned}$$

Suppose $\rho_{\omega j}^i = 0$. Then it holds that $\bar{\rho}_{\omega j}^i = 0$ and therefore $\lambda_{\omega j}^i = \bar{\lambda}_{\omega j}^i$. The equalities above hold again.

It follows that ρ^i is a best response to the prior. This contradicts that, generically, best responses are pure. $\square$

For almost every Γ and p , the set $\mathcal{S}(\Gamma, p)$ consists of a finite number of arcs and loops. Although it is not necessarily the case that these arcs and loops are smooth, the number of nondifferentiabilities is finite at most. Theorem 5.2 implies that all arcs in $\mathcal{S}(\Gamma, p)$ start and end in $\{0, 1\} \times \Sigma$. Each such path consists of a finite sequence of smooth arcs of the sets $\mathcal{S}(\Gamma, p, B^*)$. A loop in $\mathcal{S}(\Gamma, p)$ consists either of a finite sequence (at least two) of differentiable arcs of the sets $\mathcal{S}(\Gamma, p, B^*)$ or is a loop of a set $\mathcal{S}(\Gamma, p, B^*)$. See Fig. 3 for an impression of the structure of $\mathcal{S}(\Gamma, p)$.

Generically, each player i has a unique best response to the prior for all possible initial states, so generically there is exactly one point of $\mathcal{S}(\Gamma, p)$ that belongs to $\{0\} \times \Sigma$. This point is both a corner point of $\{0\} \times \Sigma$ and a boundary point of $\mathcal{S}(\Gamma, p, B^*)$, where B^* consists of the following nz strategies: for each possible initial state and for each player the best response to the prior. Given some initial state ω , the determination of the best response against the prior in state ω , involves the determination of the behavior in all other states $\bar{\omega}$ as well. So the nz independent strategies are found by solving n optimization problems, as opposed to nz independent optimization problems. By Theorem 5.3 this point is the starting point of a uniquely defined arc of $\mathcal{S}(\Gamma, p)$. This arc is the unique feasible path of $\mathcal{S}(\Gamma, p)$ that transforms prior beliefs into (stationary) equilibrium beliefs.

Theorem 5.4. *There exists an open set of stochastic games and priors $(\Gamma, p) \in \mathcal{G}(\eta, \pi, \delta) \times \Sigma$ with full Lebesgue measure such that the stochastic tracing procedure is well defined.*

By following the feasible path starting in the unique point $\mathcal{S}(\Gamma, p) \cap (\{0\} \cap \Sigma)$ we find a stationary equilibrium of the stochastic game Γ . The set $\mathcal{S}(\Gamma, p) \cap (\{1\} \times \Sigma)$ consists of all stationary equilibria of the stochastic game Γ . Precisely one of these stationary equilibria is an element of the feasible path of $\mathcal{S}(\Gamma, p)$. Any other stationary equilibrium is a boundary point of $\mathcal{S}(\Gamma, p)$ and is therefore part of some arc of $\mathcal{S}(\Gamma, p)$. A moment of reflection makes clear that the remaining stationary equilibria are pairwise connected by arcs from $\mathcal{S}(\Gamma, p)$, and so the number of stationary equilibria is odd.

Fig. 3. The structure of $\mathcal{S}(\Gamma, p)$.

Theorem 5.5. *There exists an open set of stochastic games $\Gamma \in \mathcal{G}(\eta, \pi, \delta)$ with full Lebesgue measure such that the number of stationary equilibria is odd.*

The generic oddness of the number of Nash equilibria in normal form games is a well-known result of [9,29,39]. The generic finiteness of the number of stationary equilibria in stochastic games is a recent result of Haller and Lagunoff [8]. Theorem 5.5 shows that Haller and Lagunoff's result can be sharpened to oddness.

The observations made so far suggest the following algorithm for the computation of the stationary equilibrium selected by the stochastic tracing procedure in n -person stochastic games. Define the admissible set B^* that contains for each player i the best response to the prior for all possible initial states, and start with a point $(0, \rho)$ in $\mathcal{S}(\Gamma, p, B^*)$ such that ρ_{ω}^i is a best response of player i to the prior when ω is the initial state. Theorem 5.3 implies that B^* and ρ are uniquely determined. Equalities (1)–(4) belonging to this B^* determine at least a part of the feasible path. As long as inequalities (5)–(8) hold with strict inequality we do not change our B^* . As soon as one of the inequalities from (5) or (6) gets binding, we change B^* . When the binding inequality belongs to (5), say $\rho_{\omega j}^i = 0$ while $s_{\omega j}^i \in B^*$, we delete $s_{\omega j}^i$ out of B^* . Obviously, this cannot happen for the starting B^* . If the binding inequality belongs to (6), say $\lambda_{\omega j}^i = 0$ while $s_{\omega j}^i \notin B^*$, we have to add strategy $s_{\omega j}^i$ to B^* . In both situations there is a strategy $s_{\omega j}^i$ for which $\rho_{\omega j}^i = 0$ and $\lambda_{\omega j}^i = 0$. In general this leads to a kink in the feasible path of the stochastic tracing procedure. This procedure is repeated over and over again, until inequality (8) is binding and a stationary equilibrium is found. Note that inequality (7) is only binding in the starting point.

6. Smoothing the stochastic tracing procedure

The previous section presents a method that can be used for the computation of a stationary equilibrium. A potential drawback of that method is that one has to check all the time whether the system of equations used is still appropriate by consulting the inequalities, and if not, one has to switch to a different system of equations. This

switching can be a serious problem in terms of computing time. There are $\prod_{(i,\omega) \in N \times \Omega} (2^{m_i^i} - 1)$ different sets B^* , whereas each one of them may be generated several times in the course of the algorithm.

We follow a suggestion in Garcia and Zangwill [7], also used in Herings and Peeters [14], and formulate one, everywhere differentiable, homotopy by using a well-chosen transformation of variables. Define, for $\alpha \in \mathbb{R}^{m^*}$,

$$\rho_{\omega j}^i(\alpha) = [\max\{0, \alpha_{\omega j}^i\}]^2 \quad \text{and} \quad \lambda_{\omega j}^i(\alpha) = [\max\{0, -\alpha_{\omega j}^i\}]^2.$$

After this transformation of variables, the system of equalities and inequalities becomes

$$\begin{aligned} & \tilde{u}^i(t; \omega, \rho_{\omega}^{-i}(\alpha), s_{\omega j}^i) + \delta \cdot \sum_{\bar{\omega} \in \Omega} \tilde{\pi}^i(t; \bar{\omega} \mid \omega, \rho_{\omega}^{-i}(\alpha), s_{\omega j}^i) \mu_{\bar{\omega}}^i \\ & + \lambda_{\omega j}^i(\alpha) - \mu_{\omega}^i = 0 \quad (s_{\omega j}^i \in S_{\omega}^i, \omega \in \Omega, i \in N), \\ & \lambda_{\omega j}^i(\alpha) \geq 0, \quad \rho_{\omega j}^i(\alpha) \geq 0, \quad \lambda_{\omega j}^i(\alpha) \rho_{\omega j}^i(\alpha) = 0 \\ & (s_{\omega j}^i \in S_{\omega}^i, \omega \in \Omega, i \in N), \\ & \sum_{s_{\omega j}^i \in S_{\omega}^i} \rho_{\omega j}^i(\alpha) - 1 = 0 \quad (\omega \in \Omega, i \in N). \end{aligned}$$

Owing to the proposed transformation of variables, the conditions $\lambda_{\omega j}^i(\alpha) \geq 0$, $\rho_{\omega j}^i(\alpha) \geq 0$ and $\lambda_{\omega j}^i(\alpha) \rho_{\omega j}^i(\alpha) = 0$ are trivially satisfied. We can reformulate (in)equalities in (1)–(8) that characterize the set $\mathcal{O}(\Gamma, p, B^*)$ by considering solutions $(t, \alpha, \mu) \in \mathbb{R} \times \mathbb{R}^{m^*} \times \mathbb{R}^{nz}$, with $\alpha_{\omega j}^i \geq 0$ if $s_{\omega j}^i \in B^*$ and $\alpha_{\omega j}^i \leq 0$ if $s_{\omega j}^i \notin B^*$, to

$$\begin{aligned} \text{(a)} \quad & \tilde{u}^i(t; \omega, \rho_{\omega}^{-i}(\alpha), s_{\omega j}^i) + \delta \cdot \sum_{\bar{\omega} \in \Omega} \tilde{\pi}^i(t; \bar{\omega} \mid \omega, \rho_{\omega}^{-i}(\alpha), s_{\omega j}^i) \mu_{\bar{\omega}}^i \\ & + \lambda_{\omega j}^i(\alpha) - \mu_{\omega}^i = 0 \quad (s_{\omega j}^i \in S_{\omega}^i, \omega \in \Omega, i \in N), \\ \text{(b)} \quad & \sum_{s_{\omega j}^i \in S_{\omega}^i} \rho_{\omega j}^i(\alpha) - 1 = 0 \quad (\omega \in \Omega, i \in N), \\ \text{(c)} \quad & t \geq 0, \\ \text{(d)} \quad & -t + 1 \geq 0. \end{aligned}$$

Notice that the same system of equalities and inequalities can be used, irrespective of the set B^* . The role of B^* is taken over by the sign-combinations of the components of the vector α .

Counting equations and unknowns in system (a)–(d) shows that there is one degree of freedom, and therefore one expects a 1-dimensional solution set. Consider any solution (t, α, μ) to (a)–(d). When $\alpha_{\omega j}^i = 0$, then both $\rho_{\omega j}^i(\alpha)$ and $\lambda_{\omega j}^i(\alpha)$ are zero. This implies that there are exactly two admissible subsets of S^* for which the set of (in)equalities (1)–(8) are satisfied. If along a solution curve $\alpha_{\omega j}^i$ is increasing while passing zero, then $\rho_{\omega j}^i(\alpha)$ gets positive and $B_{\text{new}}^* = B_{\text{old}}^* \cup \{s_{\omega j}^i\}$. If $\alpha_{\omega j}^i$ is decreasing

Fig. 4. The structure of $\mathcal{H}^{-1}(\{0\})$.

while passing zero, then $\lambda_{\omega j}^i(\alpha)$ gets negative and $B_{\text{new}}^* = B_{\text{old}}^* \setminus \{s_{\omega j}^i\}$. When $\alpha_{\omega j}^i$ passes zero, a kink appears in the method proposed in Section 5.

The left-hand sides of equalities (a)–(b) specify the homotopy $\mathcal{H} : [0, 1] \times \mathbb{R}^{m^*} \times \mathbb{R}^{nz} \rightarrow \mathbb{R}^{m^*} \times \mathbb{R}^{nz}$,

$$\mathcal{H}(t, \alpha, \mu) = \begin{pmatrix} \tilde{u}^i(t; \omega, \rho_{\omega}^{-i}(\alpha), s_{\omega j}^i) + \delta \cdot \sum_{\bar{\omega} \in \Omega} \tilde{\pi}^i(t; \bar{\omega} \mid \omega, \rho_{\omega}^{-i}(\alpha), s_{\omega j}^i) \mu_{\bar{\omega}}^i \\ + \lambda_{\omega j}^i(\alpha) - \mu_{\omega}^i = 0 \quad (s_{\omega j}^i \in S_{\omega}^i, \omega \in \Omega, i \in N) \\ \sum_{s_{\omega j}^i \in S_{\omega}^i} \rho_{\omega j}^i(\alpha) - 1 = 0 \quad (\omega \in \Omega, i \in N) \end{pmatrix}.$$

The homotopy function $\mathcal{H}$ is continuously differentiable.⁸ The inequalities (c) and (d) are satisfied as the homotopy takes $[0, 1]$ as the domain for the variable t . Further, H has the salient feature that its zeros describe the stochastic tracing procedure, $(t, \alpha, \mu) \in \mathcal{H}^{-1}(\{0\})$ if and only if $(t, \rho(\alpha)) \in \mathcal{S}(\Gamma, p)$.

The set $\mathcal{H}^{-1}(\{0\})$ consists of finitely many differentiable arcs and loops. All arcs start and end in $\{0, 1\} \times \mathbb{R}^{m^*} \times \mathbb{R}^{nz}$. Loops have no points in common with $\{0, 1\} \times \mathbb{R}^{m^*} \times \mathbb{R}^{nz}$. There is exactly one arc that starts in $\{0\} \times \mathbb{R}^{m^*} \times \mathbb{R}^{nz}$ and that ends in $\{1\} \times \mathbb{R}^{m^*} \times \mathbb{R}^{nz}$. All other arcs start and end in $\{1\} \times \mathbb{R}^{m^*} \times \mathbb{R}^{nz}$ and connect two points inducing stationary equilibria of the stochastic game Γ . 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 stationary equilibrium $\rho(\tilde{\alpha})$ of Γ selected by the stochastic tracing procedure. See Fig. 4 for an impression of the structure of $\mathcal{H}^{-1}(\{0\})$.

The structure of $\mathcal{H}^{-1}(\{0\})$ is even simpler than the one of $\mathcal{S}(\Gamma, p)$. Not only are complications like bifurcations, spirals, higher dimensional solutions sets, diverging

⁸ If in the variable transformation instead of the power 2 the power $k+1$ was used, k -times differentiability would have been obtained.

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 many easily implementable numerical methods, including methods to solve ordinary differential equations.

7. Implementation

The stochastic game Γ^0 naturally decomposes into n mutually independent and separate Markov decision problems, one for each player. We have shown that generically a Markov decision problem has a unique optimal pure stationary strategy. The combination of all optimal strategies (for each player one) induces the starting point of our algorithm. This point can be determined analytically since there are finitely many pure stationary strategies in each decision problem. For each player only the total discounted payoffs for each of his pure stationary strategies have to be computed and it has to be checked which strategy generates the highest payoff.

Once the starting point is determined, the numerical process starts to follow the homotopy path from that point on. We have programmed a number of Fortran-subroutines belonging to the software-package Hompack,⁹ a Fortran77 program (see [37]). Hompack provides three qualitatively different algorithms for tracking the zero curve of the homotopy: ordinary differential equation-based, normal flow, and augmented Jacobian matrix. In general, the first algorithm is the most robust of the three algorithms, but is also the slowest, sometimes by a wide margin. Being risk-averse we used this algorithm, called fixpdf, for implementation. fixpdf is an ordinary differential equation-based algorithm working with dense Jacobian matrices.

The homotopy path is parameterized by pathlength τ . Thus $t = t(\tau)$, $\alpha = \alpha(\tau)$ and $\mu = \mu(\tau)$ along the homotopy path, and $\mathcal{H}(t(\tau), \alpha(\tau), \mu(\tau)) = 0$ identically in τ . The differential equation is characterized by

$$\frac{d}{d\tau} \mathcal{H}(t(\tau), \alpha(\tau), \mu(\tau)) = \partial \mathcal{H}(t(\tau), \alpha(\tau), \mu(\tau)) \cdot \begin{pmatrix} dt/d\tau \\ d\alpha/d\tau \\ d\mu/d\tau \end{pmatrix} = 0,$$

$$\left\| \begin{pmatrix} \frac{dt}{d\tau}, \frac{d\alpha}{d\tau}, \frac{d\mu}{d\tau} \end{pmatrix} \right\|_2 = 1,$$

and initial conditions given by

$$(t(0), \alpha(0), \mu(0)) = \mathcal{H}^{-1}(\{0\}) \cap (\{0\} \times \mathbb{R}^{m^*} \times \mathbb{R}^{nz}).$$

It follows that $t(0) = 0$ and $(\alpha(0), \mu(0))$ corresponds to the best response to the prior. The homotopy path is given by the trajectory of this initial value problem. When

⁹<http://www.netlib.org/hompack/index.html>.

$t(\bar{\tau}) = 1$, $(\alpha(\bar{\tau}), \mu(\bar{\tau}))$ is a zeropoint of $\mathcal{H}$ and $\rho(\alpha(\bar{\tau}))$ is a stationary equilibrium of the stochastic game Γ with equilibrium payoff $\mu(\bar{\tau})$. Complete details for solving the initial value problem are given in [36,38].

The homotopy function that we implemented is the one as described in Section 6. We consider stochastic games with up to $z = 5$ states, $n = 5$ players, and $m = 5$ actions for each player in each state. The homotopy path to be followed is a one-dimensional object in the real Euclidean space of dimension $z \times n \times (m + 1) + 1$. It may be possible to improve on computing times by rescaling utilities and probabilities, or by using an α -transformation with a power different from 2 (but greater than 1, to keep differentiability) or multiplied by a positive scalar different from 1. An extensive digress on the optimal numerical implementation is beyond the aim of the current paper.

Within each class of stochastic games characterized by a specific choice for m , n , and z , we have generated five randomly chosen stochastic games and five randomly chosen priors. All payoffs, transition probabilities, and priors are chosen independently from one another out of the uniform distribution on $[0, 1]$. The transition probabilities and the priors are normalized to sum up to 1. The discount factor is fixed at 0.95.

Table 1 reports the average as well as the standard deviation of the time needed to compute a stationary equilibrium in a randomly generated game within the class of stochastic games determined by a specific choice for m , n and z . The computations

Table 1
Computation times in seconds

m/n	2	3	4	5
$z = 2$				
2	0.53 (0.11)	1.89 (0.99)	15.91 (13.46)	97.41 (84.32)
3	2.03 (0.62)	20.85 (6.97)	288.94 (129.96)	
4	6.54 (1.58)	146.58 (55.01)		
5	13.48 (3.31)			
$z = 3$				
2	1.40 (0.19)	6.20 (1.72)	52.19 (11.40)	334.11 (238.43)
3	4.61 (1.35)	63.14 (29.10)	1319.33 (289.82)	
4	21.20 (11.44)	1391.43 (737.45)		
5	56.62 (21.71)			
$z = 4$				
2	2.50 (0.74)	14.30 (4.85)	195.52 (203.70)	
3	14.95 (5.90)	431.57 (144.94)		
4	75.81 (41.45)			
$z = 5$				
2	3.42 (0.85)	36.64 (12.40)	435.28 (211.68)	
3	40.95 (33.00)	621.40 (258.84)		
4	313.99 (292.18)			

were performed on a PC operating under Windows 98 with an Intel Pentium II 350 MHz processor and 64 Mb workspace. The computations are executed with an end inaccuracy of 10^{-6} . The path is followed within a range of inaccuracy of 10^{-4} except for the triples $(z, n, m) \in \{(3, 3, 4), (4, 3, 3)\}$ where a maximal inaccuracy of 10^{-5} was needed in order to obtain convergence.

The average computation time increases most strongly in the number of players, followed by the number of pure actions per player per state, and finally the number of states. For instance, for the case with two pure actions per player per state and two states, an increase in the number of players from 2 to 5 leads to an increase in average computation time from 0.53 to 97.41. Fixing the number of players and the number of states, but increasing m from 2 to 5 leads to an increase in computation time from 0.53 to 13.48. Finally, when m and n equal 2, but z is increased from 2 to 5, computation time increases from 0.53 to 3.42.

The strong dependence of computation time on the number of players should not come as a surprise. Since the total number of equations to be solved equals $z \times n \times (m + 1)$, of which $z \times n \times m$ have degree $n + 1$ and the remaining $z \times n$ have degree 1, an increase in the number of players does not only give rise to more equations, but also to more complicated equations. Computation time depends also on the time it takes to evaluate the expected payoffs of all the players, which in turn is related to the total number of pure action combinations. The latter number is given by $z \times m^n$, so is linear in z and polynomial of degree n in m . It should therefore also not come as a surprise that computation time increases faster in m than in z .

Acknowledgments

We thank Anton Stoorvogel, Frank Thuijsman, and Jan van der Wal for very helpful comments.

Appendix A. Proofs

To make the proofs as transparent as possible, we need some notations and definitions from the theory of regular constraint sets as presented in [16], see also [11] for the first application of this theory in 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 $\phi: U \rightarrow V$, 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 $\phi(\bar{x}) = 0$ and $\phi(U \cap M)$ equals $\{y \in V \mid y_i = 0, i = 1, \dots, m - k, \text{ and } y_i \geq 0, i = m - k + 1, \dots, m - k + \ell(\bar{x})\}$. 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 a $(k - 1)$ -dimensional manifold, called the boundary of M .

One way to show that a set is an MGB is by showing that it is a regular constraint set. Let J^1 and J^2 be two finite index sets and let g_j for all $j \in J^1$ and h_j for all $j \in J^2$, be C^r functions defined on some open subset X of $\mathbb{R}^m$. We define

$$M[g, h] = \{x \in X \mid g_j(x) = 0, \forall j \in J^1, \text{ and } h_j(x) \geq 0, \forall j \in J^2\}.$$

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

$$\{\partial_{\bar{x}} g_j(\bar{x}) \mid j \in J^1\} \cup \{\partial_{\bar{x}} h_j(\bar{x}) \mid j \in J^0(\bar{x})\}$$

is a set of independent vectors, then $M[g, h]$ is called a C^r regular constraint set (RCS). In [16] it is shown that every C^r RCS is an $(m - |J^1|)$ -dimensional C^r MGB with $\ell(\bar{x}) = |J^0(\bar{x})|$ for every $\bar{x} \in M[g, h]$.

In this entire appendix, let a size vector η , a transition mapping π , a discount factor δ , and a prior $p \in \Sigma$ be fixed. For any tuple of utility functions u and any admissible subset B^* , define the functions $g^{B^*, u} : \mathbb{R}^{1+2m^*+nz} \rightarrow \mathbb{R}^{2m^*+nz}$ and $h^{B^*, u} : \mathbb{R}^{1+2m^*+nz} \rightarrow \mathbb{R}^{m^*+2}$ in such a way that $g^{B^*, u}$ equals the left-hand side of equalities (1)–(4) and $h^{B^*, u}$ the left-hand side of inequalities (5)–(8),

$$g^{B^*, u}(t, \rho, \lambda, \mu) = \begin{pmatrix} \tilde{u}^i(t; \omega, \rho_{\omega}^{-i}, s_{\omega j}^i) + \delta \cdot \sum_{\bar{\omega} \in \Omega} \tilde{\pi}^i(t; \bar{\omega} \mid \omega, \rho_{\omega}^{-i}, s_{\omega j}^i) \mu_{\bar{\omega}}^i \\ + \lambda_{\omega j}^i - \mu_{\omega}^i \quad (s_{\omega j}^i \in S_{\omega}^i, \omega \in \Omega, i \in N) \\ \rho_{\omega j}^i \quad (s_{\omega j}^i \notin B_{\omega}^i, \omega \in \Omega, i \in N) \\ \lambda_{\omega j}^i \quad (s_{\omega j}^i \in B_{\omega}^i, \omega \in \Omega, i \in N) \\ \sum_{s_{\omega j}^i \in S_{\omega}^i} \rho_{\omega j}^i - 1 \quad (\omega \in \Omega, i \in N) \end{pmatrix}$$

and

$$h^{B^*, u}(t, \rho, \lambda, \mu) = \begin{pmatrix} \rho_{\omega j}^i \quad (s_{\omega j}^i \in B_{\omega}^i, \omega \in \Omega, i \in N) \\ \lambda_{\omega j}^i \quad (s_{\omega j}^i \notin B_{\omega}^i, \omega \in \Omega, i \in N) \\ t \\ -t + 1 \end{pmatrix}.$$

A point $(t, \rho, \lambda, \mu) \in \mathcal{O}(\Gamma, p, B^*)$ is a solution of (1)–(8) if and only if $g^{B^*, u}(t, \rho, \lambda, \mu) = 0$ and $h^{B^*, u}(t, \rho, \lambda, \mu) \geq 0$.

Let the functions $g^{B^*} : \mathbb{R}^{1+2m^*+nz} \times \mathbb{R}^{mn} \rightarrow \mathbb{R}^{2m^*+nz}$ and $h^{B^*} : \mathbb{R}^{1+2m^*+nz} \times \mathbb{R}^{mn} \rightarrow \mathbb{R}^{m^*+2}$ be defined such that $g^{B^*}(t, \rho, \lambda, \mu, u) = g^{B^*, u}(t, \rho, \lambda, \mu)$ and $h^{B^*}(t, \rho, \lambda, \mu, u) = h^{B^*, u}(t, \rho, \lambda, \mu)$ for all $u \in \mathbb{R}^{mn}$. Fig. 5 presents the matrix of derivatives of the functions g^{B^*} and h^{B^*} with respect to all variables, where

$$E((i', \omega'), s_{\omega j}^i) = \begin{cases} 1 & \text{if } (i', \omega') = (i, \omega), \\ 0 & \text{otherwise,} \end{cases}$$

		∂_t	∂_ρ	∂_λ	∂_μ	∂_u	
		B^*	$S^* \setminus B^*$	B^*	$S^* \setminus B^*$		B^*
(1)	m^*	★	★	I	0	★	$\partial_u \tilde{u}(\cdot)$
				0	I		$S^* \setminus B^*$
(2)	m^*	0	0	I	0	0	
(3)	m^*	0	0	I	0	0	
(4)	nz	0	E	0	0	0	
(5)	m^*	0	I	0	0	0	
(6)	m^*	0	0	0	I	0	
(7)	1	1	0	0	0	0	
(8)	1	-1	0	0	0	0	
		1	m^*	m^*	nz	mn	

Fig. 5. The matrix of partial derivatives of g^{B^*} and h^{B^*} .

$$\partial_{u^{i'}(\omega', s_{\omega'}^{-i'}, s_{\omega'j'}^{i'})} \tilde{u}^i(t; \omega, \rho_\omega^{-i}, s_{\omega j}^i) = 0 \quad \text{if } s_{\omega j}^i \neq s_{\omega'j'}^{i'}, \text{ for all } s_{\omega'}^{-i'} \in S_{\omega'}^{-i'},$$

$$\sum_{s_{\omega}^{-i} \in S_{\omega}^{-i}} \partial_{u^i(\omega, s_{\omega}^{-i}, s_{\omega j}^i)} \tilde{u}^i(t; \omega, \rho_\omega^{-i}, s_{\omega j}^i) = 1.$$

The stars (★) in the matrix need not to be specified for our analysis.

Each row in Fig. 5 corresponds to one of the equalities and inequalities in (1)–(8). To make the figure more clear, derivatives with respect to ρ are first taken for $s_{\omega j}^i \in B^*$. The same applies to the derivatives with respect to λ and the ordering of the equalities in (1). From the properties of the matrix, it follows immediately that the matrices $\partial_u \tilde{u}$ and E have full row rank.

The structure of the proof of Theorem 5.2 is as follows. First, it is shown that for almost all games u , $M[g^{B^*,u}, h^{B^*,u}]$ is a regular constraint set, from which the manifold structure of $\mathcal{L}(\Gamma, p, B^*)$ follows immediately. Next it is shown that there is an open set of full measure for which the manifold structure holds. For the first part of the proof, we show that for almost all games u , for every $\bar{x} \in M[g^{B^*,u}, h^{B^*,u}]$,

$$\{\partial_x g_j^{B^*,u}(\bar{x}) \mid j \in J^1\} \cup \{\partial_x h_j^{B^*,u}(\bar{x}) \mid j \in J^0(\bar{x})\}$$

is a set of independent vectors. To show this we need three lemmas, Lemmas A.1–A.3. Lemma A.1 handles points $\bar{x}$ for which $J^0(\bar{x}) = \emptyset$, Lemma A.2 deals with points $\bar{x}$ such that the cardinality of $J^0(\bar{x})$ is one, and Lemma A.3 implies that $J^0(\bar{x})$ contains one element at most.

Lemma A.1. *Let a prior $p \in \Sigma$ and an admissible subset B^* of S^* be given. Then, for almost all u , $g^{B^*,u} \ninter \{0\}$.¹⁰*

¹⁰ Here, the symbol ' $\ninter$ ' denotes transversal intersection.

Proof. Consider a point $(\bar{t}, \bar{\rho}, \bar{\lambda}, \bar{\mu}, \bar{u})$ such that $g^{B^*}(\bar{t}, \bar{\rho}, \bar{\lambda}, \bar{\mu}, \bar{u}) = 0$. The matrix of partial derivatives of g^{B^*} at $(\bar{t}, \bar{\rho}, \bar{\lambda}, \bar{\mu}, \bar{u})$ is given by the rows (1)–(4) in Fig. 5. We show first that this matrix has full row rank, from which it follows that $g^{B^*} \not\equiv \{0\}$.

Since $\partial_u \tilde{u}(\cdot)$ has full row rank and the derivative with respect to u in (2)–(4) is zero, it is sufficient to show that the part of the matrix given by (2)–(4) has full row rank. Since the derivative with respect to λ in (3) has full row rank, whereas the derivative with respect to λ in (2) and (4) is zero, it is sufficient to show that the rows in (2) and (4) are independent. The admissibility of B^* implies that the derivative with respect to ρ_{B^*} in (4) has full row rank. Since, the derivative with respect to ρ_{B^*} in (2) is zero, the only thing left to prove is that (2) has independent rows, which is obvious from the derivative with respect to $\rho_{S^* \setminus B^*}$. Consequently, $g^{B^*} \not\equiv \{0\}$.

By the transversality theorem (see, for example, Mas-Colell [19, Theorem I.2.2]) and since g^{B^*} is a C^∞ function, it follows that the complement of the set $\{u \in \mathbb{R}^{mm} \mid g^{B^*,u} \not\equiv \{0\}\}$ has Lebesgue measure zero. $\square$

Lemma A.2. *Let a prior $p \in \Sigma$ and an admissible subset B^* of S^* be given. Moreover, let an inequality $j' \in \{1, \dots, m^* + 2\}$ be given. Then, for almost all u , $(g^{B^*,u}, h_{j'}^{B^*,u}) \not\equiv \{0\}$.*

Proof. Consider a point $(\bar{t}, \bar{\rho}, \bar{\lambda}, \bar{\mu}, \bar{u})$ such that $g^{B^*}(\bar{t}, \bar{\rho}, \bar{\lambda}, \bar{\mu}, \bar{u}) = 0$ and $h_{j'}^{B^*}(\bar{t}, \bar{\rho}, \bar{\lambda}, \bar{\mu}, \bar{u}) = 0$. The matrix of partial derivatives of $(g^{B^*}, h_{j'}^{B^*})$ at $(\bar{t}, \bar{\rho}, \bar{\lambda}, \bar{\mu}, \bar{u})$ is given in Fig. 5 by rows (1)–(4) and a single row related to $h_{j'}^{B^*}$ in (5)–(8). We show first that this matrix has full row rank, from which it follows that $(g^{B^*}, h_{j'}^{B^*}) \not\equiv \{0\}$.

If row j' belongs to (6), (7) or (8), then it follows from the derivative with respect to $\lambda_{S^* \setminus B^*}$ or from the derivative with respect to t that row j' has rank 1. Since all other derivatives in row j' are zero, it follows as in the proof of Lemma A.1 that the rows of (1)–(4) together with row j' are independent. Consider the case where row j' belongs to (5). Following the first part of the proof of Lemma A.1, it suffices to prove that (2) and (4) together with row j' are independent. Inequality $h_{j'}^{B^*}$ states that $\rho_{\omega j'}^i \geq 0$, $s_{\omega j'}^i \in B^*$, and this inequality is now required to hold with equality. Since $\sum_{s_{\omega j'}^i \in S_{\omega}^i} \rho_{\omega j'}^i - 1 = 0$, it follows that $|B_{\omega}^i| = |B^* \cap S_{\omega}^i| > 1$. Then the row related to $s_{\omega j'}^i$ together with (2) and (4) trivially form an independent set. Consequently, $(g^{B^*}, h_{j'}^{B^*}) \not\equiv \{0\}$. By the transversality theorem it follows that the complement of $\{u \in \mathbb{R}^{mm} \mid (g^{B^*,u}, h_{j'}^{B^*,u}) \not\equiv \{0\}\}$ has Lebesgue measure zero. $\square$

Lemma A.3. *Let a prior $p \in \Sigma$ and an admissible subset B^* of S^* be given. Moreover, let inequalities $j', j'' \in \{1, \dots, m^* + 2\}$ with $j' \neq j''$ be given. Then, for almost all vectors u , $(g^{B^*,u}, h_{j'}^{B^*,u}, h_{j''}^{B^*,u}) \not\equiv \{0\}$.*

Proof. Consider a point $(\bar{t}, \bar{\rho}, \bar{\lambda}, \bar{\mu}, \bar{u})$ such that $g^{B^*}(\bar{t}, \bar{\rho}, \bar{\lambda}, \bar{\mu}, \bar{u}) = 0$, $h_{j'}^{B^*}(\bar{t}, \bar{\rho}, \bar{\lambda}, \bar{\mu}, \bar{u}) = 0$, and $h_{j''}^{B^*}(\bar{t}, \bar{\rho}, \bar{\lambda}, \bar{\mu}, \bar{u}) = 0$. The matrix of partial derivatives of $(g^{B^*}, h_{j'}^{B^*}, h_{j''}^{B^*})$ at $(\bar{t}, \bar{\rho}, \bar{\lambda}, \bar{\mu}, \bar{u})$ is given in Fig. 5 by rows (1)–(4) and two rows related to $h_{j'}^{B^*}$ and $h_{j''}^{B^*}$

in (5)–(8). We show first that this matrix has full row rank, from which it follows that $(g^{B^*}, h_{j'}^{B^*}, h_{j''}^{B^*}) \not\in \{0\}$.

The case where the two rows are not equal to (7) and (8) is similar to the proof of Lemma A.2. Rows (7) and (8) are not independent. However, they cannot be binding simultaneously, because then it holds that $t = 0$ and 1. Consequently, $(g^{B^*}, h_{j'}^{B^*}, h_{j''}^{B^*}) \not\in \{0\}$. It follows that the complement of the set $\{u \in \mathbb{R}^m \mid (g^{B^*,u}, h_{j'}^{B^*,u}, h_{j''}^{B^*,u}) \not\in \{0\}\}$ has Lebesgue measure zero. $\square$

References

- [1] D. Bergemann, J. Välimäki, Learning and strategic pricing, *Econometrica* 64 (1996) 1125–1149.
- [2] M. Breton, Algorithms for stochastic games, in: T.E.S. Raghavan, T.S. Ferguson, T. Parthasarathy, O.J. Vrieze (Eds.), *Stochastic Games and Related Topics: In Honor of Professor L.S. Shapley, Theory and Decision Library C: Game Theory, Mathematical Programming and Operations Research*, Vol. 7, Kluwer Academic Publishers, Dordrecht, 1991, pp. 45–57.
- [3] M. Breton, J.A. Filar, A. Haurie, T.A. Schultz, On the computation of equilibria in discounted games, in: T. Basar (Ed.), *Dynamic Games and Applications in Economics, Lecture Notes on Economics and Mathematical Systems*, Vol. 205, Springer, Berlin, 1986.
- [4] A.H. van den Elzen, A.J.J. Talman, An algorithmic approach towards the tracing procedure for bimatrix games, *Games Econ. Behav.* 28 (1999) 130–145.
- [5] J.A. Filar, T.A. Schultz, F. Thuijsman, O.J. Vrieze, Nonlinear programming and stationary equilibria in stochastic games, *Math. Prog.* 50 (1991) 227–237.
- [6] A.M. Fink, Equilibrium in a stochastic n -person game, *J. Sci. Hiroshima Univ. Series A-I* 28 (1964) 89–93.
- [7] C.B. Garcia, W.I. Zangwill, *Pathways to Solutions, Fixed Points, and Equilibria*, Prentice-Hall, Englewood Cliffs, NJ, 1981.
- [8] H. Haller, R. Lagunoff, Genericity and Markovian behavior in stochastic games, *Econometrica* 68 (2000) 1231–1248.
- [9] J.C. Harsanyi, Oddness of the number of equilibrium points, *Int. J. Game Theory* 2 (1973) 235–250.
- [10] J.C. Harsanyi, R. Selten, *A general theory of equilibrium selection in games*, MIT Press, Cambridge, MA, 1988.
- [11] P.J.J. Herings, A globally and universally stable price adjustment process, *J. Math. Econ.* 27 (1997) 163–193.
- [12] P.J.J. Herings, Two simple proofs of the feasibility of the linear tracing procedure, *Econ. Theory* 15 (2000) 485–490.
- [13] P.J.J. Herings, A.H. van den Elzen, Computation of the Nash equilibrium selected by the tracing procedure in n -person games, *Games Econ. Behav.* 38 (2001) 89–117.
- [14] P.J.J. Herings, R.J.A.P. Peeters, A differentiable homotopy to compute Nash equilibria of n -person games, *Econ. Theory* 18 (2001) 159–185.
- [15] P.J.J. Herings, R.J.A.P. Peeters, Equilibrium selection in stochastic games, METEOR Research Memorandum 01/019, Universiteit Maastricht, Maastricht, 2001, *Int. Game Theory Rev.*, forthcoming.
- [16] H.T. Jongen, P. Jonker, F. Twilt, Nonlinear optimization in $\mathbb{R}^n$, in: *Methoden und Verfahren der mathematische Physik*, Vol. 29, Peter Lang, Frankfurt, 1983.
- [17] K.L. Judd, Computational economics and economic theory: substitutes and complements?, *J. Econ. Dynam. Control* 21 (1997) 907–942.
- [18] C.E. Lemke, J.T. Howson Jr., Equilibrium points of bimatrix games, *SIAM J. Appl. Math.* 12 (1964) 413–423.

- [19] A. Mas-Colell, *The Theory of General Economic Equilibrium: A Differentiable Approach*, Cambridge University Press, Cambridge, 1985.
- [20] E. Maskin, J. Tirole, Markov perfect equilibrium, I: observable actions, *J. Econ. Theory* 100 (2001) 191–219.
- [21] R.D. McKelvey, A. McLennan, Computation of equilibria in finite games, in: H.M. Amman, D.A. Kendrick, J. Rust (Eds.), *Handbook of Computational Economics*, Vol. I, Elsevier Science, Amsterdam, 1996, pp. 87–142.
- [22] R.D. McKelvey, T.R. Palfrey, Quantal response equilibria for normal form games, *Games Econ. Behav.* 10 (1995) 6–38.
- [23] R.D. McKelvey, T.R. Palfrey, Quantal response equilibria in extensive form games, *Exper. Econ.* 1 (1998) 9–41.
- [24] A. McLennan, The expected number of Nash equilibria of a normal form game, Mimeo, 1999.
- [25] G.S. Olley, A. Pakes, The dynamics of productivity in the telecommunications equipment industry, *Econometrica* 64 (1996) 1263–1298.
- [26] A. Pakes, R. Ericson, Empirical implications of alternative models of firm dynamics, *J. Econ. Theory* 79 (1989) 1–45.
- [27] A. Pakes, P. McGuire, Stochastic algorithms, symmetric Markov perfect equilibrium, and the ‘curse’ of dimensionality, *Econometrica* 69 (2001) 1261–1282.
- [28] T. Parthasarathy, T.E.S. Raghavan, An orderfield property for stochastic games when one player controls transition probabilities, *J. Opt. Theory Appl.* 33 (1981) 375–392.
- [29] J. Rosenmüller, On a generalization of the Lemke–Howson algorithm to noncooperative n -person games, *SIAM J. Appl. Math.* 21 (1971) 73–79.
- [30] H. Scarf, The approximation of fixed points of a continuous mapping, *SIAM J. Appl. Math.* 15 (1967) 1328–1343.
- [31] T.A. Schultz, *Mathematical Programming and Stochastic Games*, Ph.D. thesis, The John Hopkins University, Baltimore, MD, 1986.
- [32] L.S. Shapley, Stochastic games, *Proc. Natl. Acad. Sci. USA* 39 (1953) 1095–1100.
- [33] M.J. Sobel, Non-cooperative stochastic games, *Ann. Math. Statist.* 42 (1971) 1930–1935.
- [34] B. von Stengel, A.H. van den Elzen, A.J.J. Talman, Computing normal form perfect equilibria for extensive two-person games, *Econometrica* 70 (2002) 693–715.
- [35] M. Takahashi, Equilibrium points of stochastic, noncooperative n -person games, *J. Sci. Hiroshima Univ. Series A-I* 28 (1964) 95–99.
- [36] L.T. Watson, A globally convergent algorithm for computing fixed points of C^2 maps, *Appl. Math. Comput.* 5 (1979) 297–311.
- [37] L.T. Watson, S.C. Billups, A.P. Morgan, HOMPAC: A suite of codes for globally convergent homotopy algorithms, *ACM Trans. Math. Software* 13 (1987) 281–310.
- [38] L.T. Watson, D. Fenner, Chow-Yorke algorithm for fixed points or zeros of C^2 maps, *ACM Trans. Math. Software* 6 (1980) 252–260.
- [39] R. Wilson, Computing equilibria of n -person games, *SIAM J. Appl. Math.* 21 (1971) 80–87.