A Theory of Regular Markov Perfect Equilibria in Dynamic Stochastic Games: Genericity, Stability, and Purification

Ulrich Doraszelski
Harvard University*

Juan F. Escobar
Stanford University†

December 6, 2007
Preliminary and incomplete. Comments very welcome.

Abstract

This paper develops a theory of regular Markov perfect equilibria in dynamic stochastic games. We show that almost all dynamic stochastic games have a finite number of locally isolated Markov perfect equilibria that are all regular. These equilibria are essential and strongly stable. Moreover, they all admit purification. *JEL* classification numbers: C73, C61, C62.

Keywords: Dynamic stochastic games, Markov perfect equilibrium, regularity, genericity, finiteness, stability, essentiality, purifiability, applied dynamic modeling, repeated games.

*Department of Economics, Harvard University, Cambridge, MA 02138, E-mail: `doraszelski@harvard.edu`
†Department of Economics, Stanford University, Stanford, CA 94305, E-mail: `jescobar@stanford.edu`

Contents

1 Introduction

Stochastic games are central to the analysis of strategic interactions among forward-looking players in dynamic environments. Dating back to Shapley (1953), they have had a long tradition in economics. Applications of dynamic stochastic games abound and range from public finance (Bernheim and Ray 1989) and political economics (Acemoglu and Robinson 2001) to industrial organization (Bergemann and Välimäki 1996). An especially well-known example is the Ericson and Pakes (1995) model of dynamic competition in an oligopolistic industry with investment, entry, and exit that has triggered a large and active literature in industrial organization (see Doraszelski and Pakes (2007) for a survey) and, most recently, has been used also in other fields such as international trade (Erdem and Tybout 2003) and finance (Goettler, Parlour, and Rajan 2005, Kadyrzhanova 2006). In all these models the equilibrium concept is that of Markov perfect equilibrium, in which the equilibrium play is restricted to depend only on payoff relevant variables.

While several results in the literature guarantee existence of Markov perfect equilibria in dynamic stochastic games (e.g., Fink 1964, Sobel 1971, Federgruen 1978, Whitt 1980, Escobar 2007), to date very little is known about the structure of the equilibrium set in these dynamic environments. This sharply contrasts with normal form games where a large body of research is devoted to characterizing the equilibrium set and the properties of its members.

Let us discuss a few concerns unanswered by the theory of dynamic games. Consider first the problem of specification robustness. Up to now, nothing is known about whether small changes in the parameters of a dynamic model grossly alter the nature of the interaction among economic agents. This concern is especially prevalent if the researcher uses modern econometric techniques (such as Aguirregabiria and Mira 2007, Bajari, Benkard, and Levin 2007, Pakes, Ostrovsky, and Berry 2006, Pesendorfer and Schmidt-Dengler 2003) to estimate the underlying primitives of a dynamic stochastic game and therefore recovers payoffs with estimation error. It is also unknown whether differentiable comparative statics are well defined in dynamic settings. More formally, the theory of dynamic games does not provide a general answer to the question of whether there is (at least locally) a continuous, differentiable function mapping games to equilibria.

The second main open question is whether there is a sensible interpretation to behavior strategy equilibria in dynamic games. After all, when playing a non degenerate behavior strategy, at each decision node, a game player is indifferent among several pure actions. It is then natural to ask why the player should randomize precisely as mandated by his equilibrium behavior strategy. For normal form games, Harsanyi (1973a) stated and proved a result –later known as Harsanyi's purification theorem– that provided an elegant answer to this criticism, namely, a mixed strategy equilibrium can be seen and interpreted as a strict equilibrium of a nearby game of incomplete information. While the problem of purification in dynamic games has received some attention recently (Bhaskar 2000, Bhaskar, Mailath, and Morris 2007), no result characterizes a general class of dynamic environments where the question of purification can be answered.

The final, but no less important, concern is the problem of formulation robustness. As it turns out, dynamic stochastic games of incomplete information are often easier to solve numerically than their complete information counterparts. Doraszelski and Satterthwaite (2007), for example, reformulate the Ericson and Pakes (1995) model of dynamic competition as an incomplete information game with the express purpose of rendering it computationally tractable using standard algorithms (Pakes and McGuire 1994, Pakes and McGuire 2001). Dynamic stochastic games of incomplete information are also more tractable econometrically, thereby making the perturbed game the natural starting point for structurally estimating underlying primitives (Aguirregabiria and Mira 2007, Bajari, Benkard, and Levin 2007, Pakes, Ostrovsky, and Berry 2006, Pesendorfer and Schmidt-Dengler 2003). Yet, to date it has been unknown whether the nature of strategic interactions among players is sensitive to the different formulations of a dynamic economic problem.

The goal of this paper is to provide an answer to these open questions by developing a theory of regular Markov perfect equilibria in discrete-time, infinite-horizon dynamic stochastic games with a finite number of states and actions. We begin by introducing a suitable regularity notion and show that all Markov perfect equilibria of almost all dynamic stochastic games are regular, meaning that regularity is a generic property of Markov perfect equilibria. As a consequence, almost all dynamic stochastic games have a finite number of Markov perfect equilibria that are locally isolated. These equilibria are essential and strongly stable and are therefore robust to misspecifications in payoffs. Moreover, they all admit purification and can therefore be obtained as limits of equilibria of dynamic games of incomplete information as random payoff fluctuations become negligible. In sum, this paper shows that several of the most important results of the by now standard theory of regular Nash equilibria in normal form games, including genericity (Harsanyi 1973b), stability (Wu and Jiang 1962, Kojima, Okada, and Shindoh 1985), and purifiability (Harsanyi 1973a), carry over from static to dynamic environments.

Our first main insight is that, holding fixed the value of future play, the strategic situation that the players face in a given state of a dynamic stochastic game is akin to a normal form game. Consequently, a Markov perfect equilibrium of a dynamic stochastic game must satisfy the equilibrium conditions of a certain reduced one-shot game. These equilibrium conditions can be used to derive a nonlinear system of equations, $f(\sigma) = 0$, that must be satisfied by any Markov perfect equilibrium σ; we say that the equilibrium σ is regular if the Jacobian matrix $\frac{\partial f}{\partial \sigma}(\sigma)$ has full rank. Our notion of regularity is closely related to that introduced by Harsanyi (1973a, 1973b) for normal form games (and, indeed, reduces to it if players fully discount the future). While other notions of regularity have been exploited in the literature on dynamic stochastic games (e.g., Haller and Lagunoff 2000), generalizing the standard notion for normal form games in the way we do is one of the main reason that we are able to obtain our genericity, stability, and purification results.

With a notion of regularity in hand, our task is twofold. First, we show that our regularity notion is applicable to a wide range of dynamic stochastic games. Second, we show that regularity is a useful intermediate, purely mathematical definition in that regular equilibria satisfy a number of interesting economic properties.

We show that regularity is a generic property of Markov perfect equilibria in the sense that all Markov perfect equilibria of almost all dynamic games are regular. Somewhat more formally, by identifying a dynamic stochastic game with its period payoffs, it is shown that the set of games having only regular Markov perfect equilibria has full Lebesgue measure.

The proof of our main genericity result builds on the seminal work of Harsanyi (1973b). While our regularity notion facilitates the proof, the dynamic nature of the economic setting under consideration introduces nonlinearities that preclude us from simply applying Harsanyi's construction. Two insights are the key to our proof. The first is that the map relating the period payoffs of a dynamic stochastic game to the payoff matrix of the reduced one-shot game that underlies our regularity notion is invertible. This property is evident if players fully discount the future and the dynamic stochastic game collapses to a set of disjoint normal form games, but it is less than obvious in the presence of nontrivial dynamics. The second observation is that in order to study the regularity of the system $f(\sigma) = 0$, we do not have to follow Harsanyi's strategy by introducing an alternative globally defined function, relating its regular points to those of f, and then applying Sard's theorem to the alternative function (Harsanyi 1973b, van Damme 1991). In a departure from the standard treatment in the literature on normal form games, defining such an alternative function seems neither possible nor necessary in our dynamic context. Instead, we study the regularity of the system by directly applying the transversality theorem—a generalization of Sard's theorem—to it.

As a corollary to our main genericity result we deduce that almost all dynamic stochastic games have a finite number of Markov perfect equilibria that are locally isolated. While this result has already been established in an important paper by Haller and Lagunoff (2000), deriving it as part of a theory of regular Markov perfect equilibria makes for a shorter and, we believe, more transparent proof. In contrast to our approach, Haller and Lagunoff (2000) exploit a notion of regularity based on the idea that once-and-for-all deviations from the prescribed equilibrium strategies cannot be profitable. While there are clearly many types of deviations one can consider in a dynamic stochastic game, our focus on one-shot deviations has the additional advantage that it permits us to generalize several other important results for normal form games, besides generic finiteness, to dynamic stochastic games.[1]

We demonstrate that regular Markov perfect equilibria are robust to slight modifications in payoffs and, more specifically, that the equilibria of a given game can be approximated by the equilibria of nearby dynamic stochastic games. To this end, we generalize two stability properties that have received considerable attention in the literature in normal form games to our dynamic setting, namely essentiality (Wu and Jiang 1962) and strong stability (Kojima, Okada, and Shindoh 1985). Loosely speaking, a Markov perfect equilibrium is essential if it can be approximated by equilibria of nearby games; it is strongly stable if it changes uniquely and continuously with slight

[1] As a by-product of developing an algorithm for computing Markov perfect equilibria, Herings and Peeters (2004) have more recently shown that, generically, the number of Markov perfect equilibria is not only finite but also odd. While not pursued here, this result could be deduced by elaborating on our arguments.

alterations of payoffs. We show that regular equilibria are strongly stable and, therefore, essential. This result in combination with our main genericity result yields the generic essentiality and strong stability of Markov perfect equilibria.[2] We, moreover, show that the map from payoffs to equilibria is locally not only continuous but also differentiable.

Our stability results provide an immediate answer to the problem of specification robustness by ensuring that slight changes in the payoffs do not grossly alter the nature of the strategic interactions among players. They, in addition, lay the foundations for differentiable comparative statics in the form of counterfactuals, an exercise widely performed in empirical applications of dynamic stochastic games.

Next we show that regular Markov perfect equilibria admit purification, thereby extending Harsanyi's (1973a) celebrated purification theorem to dynamic stochastic games. We perturb the dynamic stochastic game by assuming that, at each decision node, a player's payoffs are subject to arbitrarily small random fluctuations known to the player but unknown to his rivals. We demonstrate that regular Markov perfect equilibria of the original complete information game can be obtained as limits of strict equilibria of this perturbed game of incomplete information as payoff fluctuations become negligible.

The purification theorem provides a compelling interpretation to behavior strategy equilibria in dynamic games according to which the original games should be seen as an idealization—a limit—of nearby games with a small amount of payoff uncertainty. It also answers the problem of formulation robustness by showing that the choice between complete and incomplete information is largely one of convenience.

While determining whether an equilibrium is purifiable seems a question of utmost concern, the problem of purification in dynamic games has received only scarce and relatively recent attention. Bhaskar (2000) and Bhaskar, Mailath, and Morris (2007) provide examples of nonpurifiable equilibrium strategies prescribing behavior that depends on payoff irrelevant history; in contrast, our result shows that equilibrium strategies conditioning only on payoff relevant variables are generically purifiable. We discuss how our results apply to equilibrium strategies conditioning on payoff irrelevant history later on. At this moment we simply point out that, to the best of our knowledge, the present work is the first one studying the purifibility of equilibrium behavior in a general and abstract class of dynamic settings.[3]

The remainder of this paper is organized as follows. Section 2 sets out the model and solution concept. Section 3 provides notation and preliminary definitions. Section 4 introduces a notion of regularity and illustrates it with two examples. Section 5 states the main genericity theorem. It also discuss some implications of the main genericity result for the finiteness of the equilibrium set. Section 6 presents stability properties and Section 7 our main purification result. Section 8

[2] Maskin and Tirole (2001) have demonstrated the generic essentiality of Markov perfect equilibria in finite-horizon dynamic stochastic games. Their result does not imply, nor is it implied by, our result.

[3] Indeed, this paper seems to be the first one aiming to develop a theory of regular equilibria in dynamic models.

contains the proofs of our main genericity and purification results. Some supporting arguments have been relegated to the Appendix.

2 Model

2.1 Dynamic Stochastic Games

A dynamic stochastic game is a dynamic system that can be in different states at different times. Players can influence the evolution of the state through their actions. The goal of a player is to maximize the net present value of his stream of payoffs.

We study dynamic stochastic games with finite sets of players, states, and actions. Let I denote the set of players, S the set of states, and A_i the set of actions of player i. To simplify the exposition, we assume that A_i is state invariant.[4]

The game proceeds as follows. The dynamic system starts from an initial state $s^{t=0}$ that is randomly drawn according to the probability distribution $\bar{q}(\cdot) \in \Delta(S)$, where $\Delta(S)$ denotes the space of probability distributions in S. After observing the initial state, players choose their actions $a^{t=0} = (a_i^{t=0})_{i \in I} \in \prod_{i \in I} A_i = A$ simultaneously and independently from each other. Now two things happen, depending on the state $s^{t=0}$ and the actions $a^{t=0}$. First, player i receives a payoff $u_i(a^{t=0}, s^{t=0}) \in \mathbb{R}$, where $u_i \colon A \times S \to \mathbb{R}$. Second, the dynamic system transits from state $s^{t=0}$ to state $s^{t=1}$ according to the probability distribution $q(\cdot; a^{t=0}, s^{t=0}) \in \Delta(S)$, with $q(s^{t=1}; a^{t=0}, s^{t=0})$ being the probability that state $s^{t=1}$ is selected. At the next round, $t = 1$, after observing the current state, $s^{t=1}$, players choose their actions $a^{t=1}$. Then players receive period payoffs $u(a^{t=1}, s^{t=1})$ and the state of the dynamic system changes again. The game goes on in this way ad infinitum. Player i's total payoff is the discounted sum of his period payoffs, where the discount factor is $\delta_i \in [0, 1[$.

2.2 Markov Perfect Equilibria

We now introduce the notion of Markov perfect equilibrium. Roughly speaking, a Markov perfect equilibrium is a subgame perfect equilibrium in which strategies depend only on payoff relevant history. In the rest of the section we provide a more precise definition of our equilibrium notion and present a characterization which will be key in our analysis.

A *stationary Markov behavior strategy* (or strategy, for short) for player i is a collection of probability distributions $(\sigma_i(\cdot, s))_{s \in S}$ such that $\sigma_i(\cdot, s) \in \Delta(A_i)$ and $\sigma_i(a_i, s)$ is the probability that player i selects action a_i in state s. We denote the set of strategies for player i as $\Sigma_i = \Delta(A_i)^{|S|}$

[4]The case of state-dependent action sets can be analyzed using exactly the same arguments that we develop below. Note, however, that even if we allow for state-dependent action sets our results do not apply to alternating-move games which, as noted by Haller and Lagunoff (2000), constitute a nongeneric subset of dynamic stochastic games. We would, however, not be surprised if suitable modifications of our arguments could be applied to these models as well. Our results apply to nonstationary finite horizon dynamic games.

and define $\Sigma = \prod_{i \in I} \Sigma_i$. We further extend $u_i(\cdot, s)$ and $q(s'; \cdot, s)$ in the obvious way to allow for randomization over $A = \prod_{i \in I} A_i$.

Definition 1 *A stationary Markov behavior strategy profile $\sigma = (\sigma_i)_{i \in I}$ is a* Markov perfect equilibrium *(or equilibrium, for short) if it is a subgame perfect equilibrium.*

We denote the set of Markov perfect equilibria of the dynamic stochastic game U by $\mathbf{Equil}(U)$. The nonemptiness of $\mathbf{Equil}(U)$ has long been established in the literature (see, e.g., Fink 1964).

We can provide a simpler characterization of equilibrium strategies using the recursive structure of the model; see Stokey and Lucas (1989) for details. A strategy profile $\sigma = (\sigma_i)_{i \in I}$ is a Markov perfect equilibrium if there exists a function $V_i \colon S \to \mathbb{R}$, for each $i \in I$, such that

1. for all $i \in I$ and all $s \in S$

$$V_i(s) = \max_{a_i \in A_i} u_i((a_i, \sigma_{-i}(\cdot, s)), s) + \delta_i \sum_{s' \in S} V_i(s') q(s'; (a_i, \sigma_{-i}(\cdot, s)), s) \tag{2.1}$$

   and

2. for all s, the profile $\sigma(\cdot, s) = (\sigma_i(\cdot, s))_{i \in I}$ is a mixed strategy Nash equilibrium of the normal form game in which each player i chooses an action $a_i \in A_i$ and, given the action profile $a \in A$, obtains a payoff

$$\mathbb{U}_i(a, s, V_i) = u_i(a, s) + \delta_i \sum_{s' \in S} V_i(s') q(ds'; a, s). \tag{2.2}$$

This characterization captures the idea that, given the continuation value functions V_i, one-shot deviations from equilibrium strategies cannot be profitable. Moreover, the observation that a Markov perfect equilibrium must prescribe Nash equilibrium behaviors for a family of normal form games will be key in the subsequent analysis.

3 Preliminaries

3.1 Notation

It will be useful to introduce some matrices that will facilitate the analysis and presentation. Consider action profiles $a^1, \ldots, a^{|A|}$ such that

$$A = \{a^1, \ldots, a^{|A|}\}.$$

Following Haller and Lagunoff (2000), we define the transition matrix $Q \in \mathbb{R}^{|S||A|\times|S|}$ as

$$Q = \begin{pmatrix} q(s_1; a^1, s_1) & \dots & q(s_{|S|}; a^1, s_1) \\ q(s_1; a^2, s_1) & \dots & q(s_{|S|}; a^2, s_1) \\ \vdots & & \vdots \\ q(s_1; a^{|A|}, s_1) & \dots & q(s_{|S|}; a^{|A|}, s_1) \\ q(s_1; a^1, s_2) & \dots & q(s_{|S|}; a^1, s_2) \\ \vdots & & \vdots \\ q(s_1; a^{|A|}, s_2) & \dots & q(s_{|S|}; a^{|A|}, s_2) \\ \vdots & & \vdots \\ \vdots & & \vdots \\ q(s_1; a^1, s_{|S|}) & \dots & q(s_{|S|}; a^1, s_{|S|}) \\ \vdots & & \vdots \\ q(s_1; a^{|A|}, s_{|S|}) & \dots & q(s_{|S|}; a^{|A|}, s_{|S|}) \end{pmatrix}$$

and, for $\sigma \in \mathbb{R}^{|S|\sum_{i\in I}|A_i|}$, we define the matrix $\mathcal{P}_\sigma \in \mathbb{R}^{|S|\times|A||S|}$ as

$$\mathcal{P}_\sigma = \begin{pmatrix} \sigma(a^1, s_1) & \dots & \sigma(a^{|A|}, s_1) & 0 & \dots & & & & & \\ 0 & \dots & 0 & \sigma(a^1, s_2) & \dots & \sigma(a^{|A|}, s_2) & 0\dots & & & \\ \vdots & & & & & & & & & \vdots \\ 0 & \dots & & & & \dots & 0 & \sigma(a^1, s_{|S|}) & \dots & \sigma(a^{|A|}, s_{|S|}) \end{pmatrix}.$$

The matrix $\mathcal{P}_\sigma$ is differentiable when seen as a function of $\sigma \in \Sigma^\epsilon$. We will consider the delta Dirac strategy $\mathbb{1}_{a_i}$ putting probability 1 on action $a_i \in A_i$ in all state $s \in S$. Therefore, for any $\sigma_{-i} \in \Sigma^\epsilon_{-i}$, $(\mathbb{1}_{a_i}, \sigma_{-i})$ is well defined in Σ^ϵ and $\mathcal{P}_{(\mathbb{1}_{a_i},\sigma_{-i})}$, simply denoted $\mathcal{P}_{a_i,\sigma_{-i}}$, is a well defined matrix. By writing $\mathcal{P}^s_\sigma$ we denote the row s of $\mathcal{P}_\sigma$.

Throughout the paper, $\mathbb{I}_r$ denotes the $\mathbb{R}^{r\times r}$ identity matrix, where $r \in \mathbb{N}$.

3.2 Definition of Σ^ϵ

For our proofs, it is necessary to introduce an open set Σ^ϵ which strictly contains Σ. In particular, Σ^ϵ contains elements which are not proper strategies.

The definition of Σ^ϵ is as follows. First, note that for each $\sigma \in \Sigma$, $\mathbb{I}_{|S|} - \delta_i \mathcal{P}_\sigma Q$ and $\mathbb{I}_{|A||S|} - \delta_i Q\mathcal{P}_\sigma$ both are invertible. Indeed, $\mathcal{P}_\sigma Q$ and $Q\mathcal{P}_\sigma$ are both stochastic matrix and so $\mathbb{I}_{|S|} - \delta_i \mathcal{P}_\sigma Q$ and $\mathbb{I}_{|A||S|} - \delta_i Q\mathcal{P}_\sigma$ have strictly dominant diagonals. For each $\bar\sigma \in \Sigma$, we can find $\epsilon_{\bar\sigma} > 0$ such that for all $\sigma \in \mathbb{R}^{\sum_{i\in I}|A_i||S|}$ satisfying $\|\bar\sigma - \sigma\|_\infty < \epsilon_\sigma$ (here $\|x\|_\infty = \max_n |x_n|$), $\mathbb{I}_{|S|} - \delta_i \mathcal{P}_\sigma Q$ and $\mathbb{I}_{|A||S|} - \delta_i Q\mathcal{P}_\sigma$ are invertible. Since Σ is compact, we can take a finite covering $\left(B(\bar\sigma^j, \epsilon^j)\right)_{j\in J}$ of Σ, where $\epsilon^j = \epsilon_{\bar\sigma^j}$ and $|J| < \infty$. Define $\Sigma^\epsilon = \cup_{j\in J} B(\bar\sigma^j, \epsilon^j)$.

3.3 Continuation Values

The function $V_i \colon S \to \mathbb{R}$ in equation (2.1) is the so called *equilibrium value function* for player i. $V_i(s)$ is the discounted expected value of the stream of payoffs to player i if the dynamic system is currently in state s. That is, $V_i(s)$ is the equilibrium value of continued play to player i starting from state s.

If $\sigma \in \Sigma$ is an equilibrium, then equation (2.1) reduces to

$$V_i(s) = u_i(\sigma(\cdot, s), s) + \delta_i \sum_{s' \in S} V_i(s') q(s'; \sigma(\cdot, s), s). \tag{3.1}$$

The collection of equations (3.1) constitutes a system of $|S|$ linear equations in the $|S|$ unknown values $\big(V_i(s)\big)_{s \in S}$ and can be written as

$$V_i(\cdot) = \mathcal{P}_\sigma U_i + \delta_i \mathcal{P}_\sigma Q V_i(\cdot), \tag{3.2}$$

where $\mathcal{P}_\sigma \in \mathbb{R}^{|S| \times |A||S|}$ and $Q \in \mathbb{R}^{|S||A| \times |S|}$ are defined in Subsection 3.1. The system of equations (3.2) can be equivalently written as

$$\Big(\mathbb{I}_{|S|} - \delta_i \mathcal{P}_\sigma Q\Big) V_i(\cdot) = \mathcal{P}_\sigma U_i$$

Importantly, this system of equations is well defined not only if $\sigma \in \Sigma$ is not an equilibrium but also if $\sigma \in \Sigma^\epsilon \setminus \Sigma$ is not a strategy. Moreover, by definition of Σ^ϵ, for all $\sigma \in \Sigma^\epsilon$, the matrix $\mathbb{I}_{|S|} - \delta_i \mathcal{P}_\sigma Q$ is invertible. Therefore, the system has a unique solution given by

$$V_i(\cdot, \sigma) = \big(\mathbb{I}_{|S|} - \delta_i \mathcal{P}_\sigma Q\big)^{-1} \mathcal{P}_\sigma U_i. \tag{3.3}$$

In other words, given distributions over actions, equation (3.3) provides us a formula to pin down the value of continuation play. Since $\mathcal{P}_\sigma$ is continuously differentiable as a function of $\sigma \in \Sigma^\epsilon$, so is $V_i(\cdot, \sigma)$.

4 Regular Equilibria

4.1 Definition of Regular Equilibrium

Our notion of regularity is based on the observation that, given continuation values, the strategic situation that the players face in a given state s is akin to a normal form game. We combine this observation, formally captured in the alternative characterization of Markov perfect equilibria provided in Section 2.2, with the development of Section 3.3 to derive a set of necessary conditions for Markov perfect equilibria.

More formally, for $a \in A$, $s \in S$, and $\sigma \in \Sigma^\epsilon$, the expression

$$\mathbb{U}_i(a, s, V_i(\cdot, \sigma)) = u_i(a, s) + \delta_i \sum_{s' \in S} V_i(s', \sigma) q(ds'; a, s). \tag{4.1}$$

is the net present value of the stream of payoffs to player i if the current action profile is a, the current state is s and play in subsequent periods accords σ. This expression highlights the fact that the payoffs in the family of normal form games our regularity notion is based on is induced by continuation strategies through continuation values. Moreover, $\sigma \in \Sigma$ is a Markov perfect equilibrium if and only if for all $s \in S$ the profile $\sigma(\cdot, s) = (\sigma_i(\cdot, s))_i$ is a mixed strategy Nash equilibrium for the normal form game in which, given an action profile $a \in A$, player i's payoff is given by (4.1).

Now, given $\sigma \in \Sigma^\epsilon$, the payoff to player i from playing $a_i \in A_i$ in the one-shot game induced in state $s \in S$ is given by the function $\mathcal{U}_i^U : A_i \times S \times \Sigma^\epsilon \to \mathbb{R}$, where

$$\mathcal{U}_i^U(a_i, s, \sigma) = u_i((a_i, \sigma_{-i}(\cdot, s)), s) + \delta_i \sum_{s' \in S} V_i(s', \sigma) q(s'; (a_i, \sigma_{-i}(\cdot, s)), s),$$

where now we highlight the dependence of this payoff function on the game $U \in \mathbb{R}^{|I||A||S|}$. An equivalent way to write the fact that, when $\sigma \in \Sigma$ is an equilibrium, for no player and no state one-shot deviations can be optimal is as follows: For all $i \in I$ and all $s \in S$,

$$\sigma_i(a_i, s) > 0 \text{ implies that } a_i \in \arg\max \left\{ \mathcal{U}_i^U(a_i', s, \sigma) \mid a_i' \in A_i \right\}. \tag{4.2}$$

Consider a collection of actions $a_i^s \in A_i$ for all $i \in I$ and $s \in S$. Define the (i, a_i, s) component of the function $f \colon \Sigma^\epsilon \times \mathbb{R}^{|I||A||S|} \to \mathbb{R}^{|S| \sum_{i \in I} |A_i|}$ by

$$f_{i,a_i,s}(\sigma, U) = \begin{cases} \sigma_i(a_i, s)\Big(\mathcal{U}_i^U(a_i, s, \sigma) - \mathcal{U}_i^U(a_i^s, s, \sigma)\Big) & \text{if} \quad a_i \neq a_i^s, \\ \sum_{a_i \in A_i} \sigma_i(a_i, s) - 1 & \text{if} \quad a_i = a_i^s. \end{cases}$$

Equation (4.2) implies that if σ is an equilibrium of the game U such that $\sigma(a_i^s, s) > 0$ for all $i \in I$ and $s \in S$, then

$$f(\sigma, U) = 0.$$

We say that the equilibrium is *regular* if the Jacobian of f with respect to σ, $\frac{\partial f(\sigma, U)}{\partial \sigma}$, has full rank for some selection of actions a_i^s such that $\sigma(a_i^s, s) > 0$ for all $i \in I$ and $s \in S$. An equilibrium that is not regular is said to be *irregular*.

Our notion of regularity is reminiscent of that introduced by Harsanyi (1973a, 1973b) for normal form games. Indeed, if $\delta_i = 0$ for all $i \in I$, our notion of regularity reduces to the standard notion. But even if $\delta_i > 0$ for some $i \in I$, our notion remains closely related to the standard notion because we base it on the equilibrium conditions for a reduced one-shot game. This observation, while simple, permits us to generalize several important results for normal form games to dynamic stochastic games.

4.2 Examples

This subsection provides examples of dynamic stochastic games having regular equilibria. Our first example is a simple version of the industry dynamics model proposed by Ericson and Pakes (1995) (see also Doraszelski and Satterthwaite 2007).

Example 2 (Exit Game) Consider a market having two firms so that the set of players is $I = \{1, 2\}$. The state space is $S = \{(1, 1), (1, 0), (0, 1), (0, 0)\}$ where given $s = (s(1), s(2)) \in S$, $s(i) = 1$ means that firm i is in the market while $s(i) = 0$ means firm i is out of the market. In this model, action sets are state-dependent; when $s = (1, 1)$ firm i's action set is $\{\text{exit}, \text{stay}\}$, when $s(i) \neq s(-i) = 0$ firm i's action set is $\{\text{stay}\}$, while when $s(i) = 0$ firm i's action space is $\{\text{exit}\}$. While in the description of the model we assumed that the action spaces are state-independent, this simple example will illustrate how easy is to consider the more general setting of state-dependent action sets.

States $s \neq (1, 1)$ are absorbent, meaning that after reaching state s the state of the game remains stuck at s for ever. If the state is $s = (1, 1)$ and firm i decides to exit (resp. stay), then in the next period the state s' will be such that $s'(i) = 0$ (resp. $s'(i) = 1$).

Payoffs are as follows. If the market is a monopoly (so the state is other than $(1, 1)$, then the monopolist receives a payoff π^M. If both firms are in the market, then each firm receives a payoff π^D. If the state s is such that $s(i) = 0$, firm i does not receive any payoff. If firm i is in the market and decides to exit, then regardless of what the other firm does, i additionally receives a scrap value ϕ. Firm i's discount factor is $\delta_i = \delta_{-i} = \delta$. We assume that

$$\frac{\delta}{1-\delta}\pi^D < \phi < \frac{\delta}{1-\delta}\pi^M,$$

meaning that, while a monopoly is viable, a duopoly is not. This completes the description of our dynamic exit game.

A strategy for player i consists of distributions $\sigma_i(\cdot, s)$, for each $s \in S$, such that $\sigma_i(\text{exit}, (1, 1)) + \sigma_i(\text{stay}, (1, 1)) = 1$ and for all $s \neq (1, 1)$, $\sigma_i(a_i, s) = 1$ where a_i is the only action available to i in state s.

In this model, the only nontrivial decision is whether or not to exit when both firms are in the market. The model has three equilibria, two of them in pure strategies. The only symmetric equilibrium, denoted $\bar{\sigma}$, is fully characterized by the exit probability

$$\bar{\sigma}_i(\text{exit}, (1, 1)) = \frac{(1-\delta)\phi - \delta\pi^M}{\delta\Big(\frac{\pi^M}{1-\delta} - \pi^D + \phi\Big)}.$$

We will show that the symmetric equilibrium $\bar{\sigma}$ is regular.

Let us first construct the function $V_i(\cdot, \sigma)$ described in equation (3.3). While deriving the four dimensional vector $V_i(\cdot, \sigma)$ by using that equation is doable, here we proceed by exploiting the simple dynamic structure of the exit game. Note that if $s(i) = 0$ then $V_i(s, \sigma) = 0$, while if $s(i) = 1$ and $s(-i) = 0$, $V_i(s, \sigma) = \pi^M/(1-\delta)$. To compute $V_i((1,1), \sigma)$, note that

$$V_i((1,1),\sigma) =$$

$$\pi^D + \sigma_i(\text{exit},(1,1))\Big\{\phi + \delta 0\Big\} + \sigma_i(\text{stay},(1,1))\Big\{\delta\sigma_{-i}(\text{exit},(1,1))\frac{\pi^M}{1-\delta} + \delta\sigma_{-i}(\text{stay},(1,1))V_i(1,1)\Big\}.$$

This expression allows us to pin down $V_i((1,1), \sigma)$ as

$$V_i((1,1),\sigma) = \frac{\pi^D + \phi\sigma_i(\text{exit},(1,1)) + \frac{\delta}{1-\delta}\pi^M\sigma_i(\text{stay},(1,1))\sigma_{-i}(\text{exit},(1,1))}{1 - \delta\sigma_i(\text{stay},(1,1))\sigma_{-i}(\text{stay},(1,1))},$$

which is a differentiable function of σ.

We are now in a position to write the 10 dimensional function f, where we set $a_i^{(1,1)} = (1,1)$. If $s \neq (1,1)$, then the i, a_i, s component of f is

$$\sigma_i(a_i, s) - 1,$$

where a_i is the only action available to i in state s. Now, for the state $s = (1,1)$ we can, for each i, write the two components of f related to i and s as

$$\sigma_i(\text{exit},(1,1)) + \sigma_i(\text{stay},(1,1)) - 1$$

and

$$\sigma_i(\text{exit},(1,1))\Big\{\pi^D + \phi - \Big(\pi^D + \delta\big(\sigma_{-i}(\text{exit},(1,1))\frac{\pi^M}{1-\delta} + \sigma_{-i}(\text{stay},(1,1))V_i((1,1),\sigma)\big)\Big)\Big\}.$$

After computing the determinant of the Jacobian of f at $\bar{\sigma}$, $\frac{\partial f}{\partial \sigma}(\bar{\sigma})$, it can be seen that $\bar{\sigma}$ is regular.

Our second example shows how a repeated prisoners' dilemma and its trigger strategies can be embedded into our dynamic stochastic game setting.

Example 3 (Repeated Prisoners' Dilemma) Consider the repeated prisoners' dilemma with two players $I = \{1,2\}$ and two actions $A_i = \{C, D\}$ per player, and a common discount factor $\delta \in]0,1[$. The payoff matrix is

	C		D	
C	1	1	$-l$	$1+g$
D	$1+g$	$-l$	0	0

Let $\mathcal{H}$ be the set of histories of the repeated game. Define the strategy $\bar{a}_i \colon \mathcal{H} \to A_i$ by

$$\bar{a}_i(a_1^{t=0}, a_2^{t=0}, \ldots, a_1^{t=T}, a_2^{t=T}) = \begin{cases} D & \text{if } a_j^{t=t'} = D \text{ for some } j \text{ and } t', \\ C & \text{otherwise.} \end{cases}$$

To represent this repeated game setting within our dynamic framework, let $S = \{\text{On}, \text{Off}\}$ be a state space and define the transition function by

$$q(\text{On}; a, s) = \begin{cases} 1 & \text{if } a = (C, C) \text{ and } s = \text{On} \\ 0 & \text{otherwise} \end{cases}$$

and $\bar{q}(\text{On}) = 1$ (with $q(\text{Off}; a, s) = 1 - q(\text{On}; a, s)$) and $\bar{q}(\text{Off}) = 0$). Consider the stationary Markov behavior strategy

$$a_i(s) = \begin{cases} C & \text{if } s = \text{On}, \\ D & \text{otherwise.} \end{cases}$$

It is relatively easy to see that the repeated-game strategy constitute a subgame perfect equilibrium of the repeated game if and only if the stationary Markov behavior strategy constitute a Markov perfect equilibrium of the dynamic stochastic game. Moreover, the observed behaviors under both strategies coincide.

Now, observe that when $\delta > 1/(1+g)$, the strategies not only constitute equilibrium strategies but the incentives hold strictly. In other words, the Markov perfect equilibrium of the dynamic stochastic game is strict, as defined in Section 6.1. Proposition 6 implies that a, being a strict equilibrium, is regular.

We will return to the discussion of repeated game strategies later on in the paper.

5 Genericity of Regular Equilibria

Before establishing the properties satisfied by regular equilibria, we show that regularity is a property satisfied by all equilibria of a large set of models. To do that, we identify a dynamic stochastic game with its period payoff functions $(u_i)_{i \in I}$. We let $U_i = (u_i(a, s))_{(a \in A, s \in S)} \in \mathbb{R}^{|A||S|}$ denote the vector of payoffs of player i and $U = (U_i)_{i \in I} \in \mathbb{R}^{|I||A||S|}$ the vector of payoffs of all players. We endow the set of games with the Lebesgue measure λ and say that a property is *generic* if it does not hold at most on a closed subset of measure zero.[5] In this case we say that the property holds for *almost all games* $U \in \mathbb{R}^{|I||A||S|}$.

The following is the main result of this section.

[5] See, for example, Haller and Lagunoff (2000) for a discussion of alternative genericity notions.

Theorem 4 *For almost all games $U \in \mathbb{R}^{|I||A||S|}$, all equilibria are regular.*

The proof of this theorem, which we detail in Section 8.1, proceeds as follows. We first consider the set of dynamic stochastic games having equilibria in which some player puts zero weight on some of his best replies. The set of such games has a small dimension and so does, therefore, the subset of games having irregular equilibria. We then consider the set of games having equilibria in which all players put positive weight on all their best replies (we call these equilibria quasi-strict, see Section 6 for a formal definition). We observe that within this class we can restrict attention to games having completely mixed equilibria. For these games we show that the derivative of f with respect to the pair (σ, U) has full rank. An application of the transversality theorem—a generalization of Sard's theorem—then yields the desired result, namely that the derivative of f with respect to the strategy profile σ has full rank for almost all games U.

Our proof is similar to that offered by Harsanyi (1973b). One difference is, of course, the technical difficulties that have to be surmounted when working with dynamic stochastic games instead of normal form games. The presence of nontrivial dynamics introduces several nonlinearities that precludes us from simply applying Harsanyi's construction. Indeed, the family of induced normal form games is itself endogenous to strategies. Two insights facilitate our analysis. The first observation is that, given a state and a strategy profile to be followed from next period on, the map relating a dynamic stochastic game to the payoff matrix of the reduced one-shot game is invertible. More formally, we show that the map $U_i \mapsto U_i + \delta_i Q V_i(\cdot, \sigma)$ is linear and invertible for all $i \in I$ and $\sigma \in \Sigma^\epsilon$. This property is evident for normal form games where the term $\delta_i Q V_i(\cdot, \sigma)$ vanishes, but to establish it in our dynamic context we have to introduce suitable notation and to carry out a number of manipulations (see Section 8.1.1). The second observation facilitating our analysis is that in order to study the regularity of f, we do not have to work with an alternative globally defined function that maps strategies to payoffs, relate its regular points to those of f, and then apply Sard's theorem to it. In a departure from this approach —introduced by Harsanyi (1973b) and further illustrated by van Damme (1991)—, defining such an alternative function seems neither possible nor necessary in our dynamic context. Instead, we study the regularity of f by directly applying the transversality theorem to it (see Section 8.1.3 for details).

Any regular equilibrium is locally isolated as a consequence of the implicit function theorem. A dynamic stochastic game having equilibria that are all regular has a compact equilibrium set consisting of isolated points; therefore the equilibrium set has to be finite. We state this result in the following corollary.

Corollary 5 (Haller and Lagunoff (2000)) *For almost all games $U \in \mathbb{R}^{|I||A||S|}$, the number of equilibria is finite.*

The above result has already been established in an important paper by Haller and Lagunoff (2000). These authors exploit a notion of regularity derived from the first-order necessary conditions for an equilibrium of the dynamic stochastic game. This system of equations captures the idea that once-and-for-all deviations from the prescribed equilibrium strategies cannot be profitable. There are clearly many types of deviations one can consider in a dynamic stochastic game, and Haller and Lagunoff (2000) choose a different approach than we do. This is so partly because they are not interested in developing a theory of regular equilibria but only in proving the above finiteness result.

While Haller and Lagunoff's (2000) approach to defining a notion of regularity is interesting, we think that our focus on deviations from a reduced one-shot game has three advantages. First, it provides a simple and intuitive generalization of the standard regularity notion for normal form games. Second, it makes for a shorter and more transparent proof of the above finiteness result.[6] Third, and perhaps most important, our regularity notion permits us to generalize several other important results for normal form games to dynamic stochastic games; we would be surprised if the regularity notion implicit in Haller and Lagunoff's (2000) work could be exploited to derive the results we provide in the rest of the paper.

6 Stability Properties of Regular Equilibria

This section explores the notions of essential and strongly stable equilibria in dynamic games. Before studying these two stability properties, we introduce the notions of strict and quasi-strict equilibria. These concepts both help us to clarify the proofs below and are invoked extensively again in Section 8. At the end of the section, we discuss the implications of these results for applied work.

6.1 Strict and Quasi-Strict Equilibria

Given a strategy profile $\sigma \in \Sigma$, we define the set of (pure) *best replies* for player i in state s as

$$B_i(\sigma, s) = \arg\max_{a_i \in A_i} \mathcal{U}_i^U(a_i, s, \sigma).$$

We also define the *carrier* $C_i(\sigma, s) \subseteq A_i$ of player i in state s as the set of actions a_i with $\sigma_i(a_i, s) > 0$. We finally define $B(\sigma) = \prod_{i\in I} \prod_{s\in S} B_i(\sigma, s)$ and $C(\sigma) = \prod_{i\in I} \prod_{s\in S} C_i(\sigma, s)$.

If σ is an equilibrium, then $C_i(\sigma, s) \subseteq B_i(\sigma, s)$ for all $i \in I$ and $s \in S$. The equilibrium is *quasi-strict* if $B_i(\sigma, s) = C_i(\sigma, s)$ for all $i \in I$ and $s \in S$. This means that all players put (strictly) positive weight on all their best replies. We further say that the equilibrium is *strict* if the set of best replies is always a singleton, i.e., $|B_i(\sigma, s)| = 1$ for all $i \in I$ and $s \in S$.

A strict equilibrium is also quasi-strict, but how these concepts relate to regularity is not immediately apparent. The following result, as well as its proof, resembles a well known result for

[6] We discuss how our arguments yield simpler finiteness proof right after Lemma 14.

normal form games (see, e.g., Corollary 2.5.3 in van Damme 1991).

Proposition 6 *Every strict equilibrium is regular. Every regular equilibrium is quasi-strict.*

Proof. Define $J(\sigma) = \frac{\partial f(\sigma,U)}{\partial \sigma}$ to be the Jacobian of f with respect to σ and consider the submatrix $\bar{J}(\sigma)$ obtained from $J(\sigma)$ by crossing out all columns and rows corresponding to components (a_i, s) with $a_i \notin C_i(\sigma, s)$. For all pairs (a_i, s) with $a_i \notin C_i(\sigma, s)$ we have

$$\frac{\partial f_{i,a_i,s}(\sigma, U)}{\partial \sigma_i(a_i, s)} = \mathcal{U}_i^U(a_i, s, \sigma) - \mathcal{U}_i^U(a_i^s, s, \sigma),$$

while for all $j \in I$ and $\widetilde{a}_j \in A_j$, with $\widetilde{a}_j \neq a_i$ if $j = i$, we have

$$\frac{\partial f_{i,a_i,s}(\sigma, U)}{\partial \sigma_j(\widetilde{a}_j, s)} = 0.$$

It follows that

$$|\det (J(\sigma))| = |\det \left(\bar{J}(\sigma)\right)| \Big| \prod_{i \in I} \prod_{s \in S} \prod_{a_i \notin C_i(\sigma,s)} \left[\mathcal{U}_i^U(a_i, s, \sigma) - \mathcal{U}_i^U(a_i^s, s, \sigma)\right] \Big|. \tag{6.1}$$

If the equilibrium $\sigma \in \Sigma$ is strict, then $\{a_i^s\} = C_i(\sigma, s) = B_i(\sigma, s)$ for all $i \in I$ and $s \in S$. Therefore, $\det \left(\bar{J}(\sigma)\right) = 1$ and $\mathcal{U}_i^U(a_i, s, \sigma) - \mathcal{U}_i^U(a_i^s, s, \sigma) < 0$ for all pairs (a_i, s) with $a_i \notin C_i(\sigma, s)$. It follows that $\det \left(J(\sigma)\right) \neq 0$ so that the equilibrium is regular. On the other hand, if the equilibrium is regular, then each of the terms on the right hand side of equation (6.1) is nonzero. Hence, $\mathcal{U}_i^U(a_i, s, \sigma) < \mathcal{U}_i^U(a_i^s, s, \sigma)$ for all pairs (a_i, s) with $a_i \notin C_i(\sigma, s)$; this corresponds to the definition of quasi-strictness. ■

6.2 Essential and Strongly Stable Equilibria

We say that an equilibrium $\bar{\sigma}$ of game $\bar{U}$ is *strongly stable* if there exist neighborhoods $\mathcal{N}_{\bar{U}}$ of $\bar{U}$ and $\mathcal{N}_{\bar{\sigma}}$ of $\bar{\sigma}$ such that the mapping **equil**: $\mathcal{N}_{\bar{U}} \to \mathcal{N}_{\bar{\sigma}}$ defined by $\mathbf{equil}(U) = \mathbf{Equil}(U) \cap \mathcal{N}_{\bar{\sigma}}$ is single-valued and continuous. In words, an equilibrium is strongly stable if the equilibrium correspondence looks locally like a continuous function. This definition generalizes that introduced for normal form games by Kojima, Okada, and Shindoh (1985).

Proposition 7 *Every regular equilibrium is strongly stable.*

Proof. Let $\bar{\sigma}$ be a regular equilibrium of game $\bar{U}$. Since $\frac{\partial f(\bar{\sigma},\bar{U})}{\partial \sigma}$ has full rank, the implicit function theorem implies the existence of open neighborhoods $\mathcal{N}_{\bar{U}}$ of $\bar{U}$ and $\mathcal{N}_{\bar{\sigma}}$ of $\bar{\sigma}$ and a differentiable function $\widetilde{\sigma}\colon \mathcal{N}_{\bar{U}} \to \mathcal{N}_{\bar{\sigma}}$ such that for all $U \in \mathcal{N}_{\bar{U}}$, $\widetilde{\sigma}(U)$ is the unique solution $\sigma \in \mathcal{N}_{\bar{\sigma}}$ of $f(\sigma, U) = 0$. We can choose $\mathcal{N}_{\bar{U}}$ and $\mathcal{N}_{\bar{\sigma}}$ small enough so that for all $i \in I$, $s \in S$, and $a_i \in A_i$: (i)

If $\bar{\sigma}_i(a_i, s) > 0$, then $\sigma_i(a_i, s) > 0$ for all $\sigma \in \mathcal{N}_{\bar{\sigma}}$; (ii) If $\mathcal{U}_i^{\bar{U}}(a_i, s, \bar{\sigma}) - \mathcal{U}_i^{\bar{U}}(a_i^s, s, \bar{\sigma}) < 0$, then $\mathcal{U}_i^U(a_i, s, \sigma) - \mathcal{U}_i^U(a_i^s, s, \sigma) < 0$ for all $(\sigma, U) \in \mathcal{N}_{\bar{\sigma}} \times \mathcal{N}_{\bar{U}}$.

Denote $\widetilde{\sigma}(U)$ by σ^U. From (i) it follows that

$$\sigma_i^U(a_i, s) > 0 \text{ for all } a_i \in C_i(\bar{\sigma}, s). \tag{6.2}$$

Now, for $a_i \notin C_i(\bar{\sigma}, s)$, the fact that $\bar{\sigma}$ is regular and so quasi-strict implies that $\mathcal{U}_i^{\bar{U}}(a_i, s, \bar{\sigma}) < \mathcal{U}_i^{\bar{U}}(a_i^s, s, \bar{\sigma})$. From (ii), it follows that $\mathcal{U}_i^U(a_i, s, \bar{\sigma}) < \mathcal{U}_i^U(a_i^s, s, \bar{\sigma})$ and, by definition of σ^U, $\sigma^U(a_i, s) = 0$. It follows that for all $U \in \mathcal{N}_{\bar{U}}$

$$\mathcal{U}_i^U(a_i, s, \sigma^U) \leq \mathcal{U}_i^U(a_i^s, s, \sigma^U) \text{ with strict inequality if and only if } a_i \notin C_i(\bar{\sigma}, s). \tag{6.3}$$

Conditions (6.2) and (6.3) imply that $\sigma^U \in \mathbf{Equil}(U)$ for all $U \in \mathcal{N}_{\bar{U}}$. Moreover, σ^U is the only equilibrium of U in $\mathcal{N}_{\bar{\sigma}}$ because any other equilibrium σ would satisfy $f(\sigma, U) = 0$; as a result $\mathbf{equil}(U) = \mathbf{Equil}(U) \cap \mathcal{N}_{\bar{\sigma}} = \sigma^U$. Since σ^U is differentiable, $\mathbf{equil}(U) = \sigma^U$ is differentiable (and so continuous). ■

The idea behind the proof is to show that close enough to a regular equilibrium $\bar{\sigma}$ of $\bar{U}$, the system $f(\sigma, U) = 0$ fully characterizes the equilibrium map **Equil**. This logic mimics that of Theorem 2.5.5 in van Damme (1991). The proof also presents an argument we further exploit in the proof of our main purification result (see Section 8.2 for details). Importantly, the proof shows that the locally defined equilibrium map **equil** is differentiable, and therefore local comparative statics exercises can be carried out.

Let us now introduce the notion of essentiality. We say that an equilibrium $\bar{\sigma}$ of game $\bar{U}$ is *essential* if for every neighborhood $\mathcal{N}_{\bar{\sigma}}$ there exists a neighborhood $\mathcal{N}_{\bar{U}}$ such that for all games $U \in \mathcal{N}_{\bar{U}}$, there exists $\sigma \in \mathbf{Equil}(U) \cap \mathcal{N}_{\bar{\sigma}}$. In words, an equilibrium is essential if it can be approximated by equilibria of nearby games. Since any strongly stable equilibrium can be approximated by equilibria of nearby games, the following result is immediate.

Proposition 8 *Every strongly stable equilibrium is essential.*

Theorem 4, Proposition 7, and Proposition 8 permit us to deduce the following corollary that generalizes a well known result for normal form games due to Wu and Jiang (1962) to our dynamic context.

Corollary 9 *For almost all games $U \in \mathbb{R}^{|I||A||S|}$, all equilibria are essential.*

6.3 Discussion

The literature on normal form games has forcefully argued that an equilibrium should be stable against perturbations of the payoffs because the data of the game are usually not known exactly.

This argument is especially pertinent in our setting if the researcher uses modern econometric techniques (such as Aguirregabiria and Mira 2007, Bajari, Benkard, and Levin 2007, Pakes, Ostrovsky, and Berry 2004, Pesendorfer and Schmidt-Dengler 2003) to estimate the underlying primitives of a dynamic stochastic game and therefore recovers payoffs with estimation error. It then becomes critical that slight changes in the payoffs do not grossly alter the nature of the strategic interactions among players.

The fact that essential equilibria are the norm, rather than the exception, implies that we can be relatively confident about the robustness of the conclusions reached when employing this econometric tools. Moreover, the observation that regular equilibria are additionally strongly stable and are locally fully characterized by the equation $f(\sigma, U) = 0$ lays the foundations for comparative statics telling us that these exercises are, at least in a local sense, well defined. (As an illustration of these results, note that the symmetric regular equilibrium presented in Example 2 is differentiable in the parameters of the model.)

While the problem of computation of equilibria is beyond the scope of this work, it seems appropriate to discuss some implications of our results to the numerical implementation of the comparative statics exercises discussed above. Consider a game $\bar{U}$ and an equilibrium $\bar{\sigma}$. In applied work, the fundamentals of the game $\bar{U}$ and the equilibrium behavior $\bar{\sigma}$ could be derived from data by using the econometric methods currently available. Suppose now we are interested in knowing how the equilibrium looks like when we slightly alter the payoff structure; say the new game is characterized by $\hat{U}$. Our results suggest that assuming the existence of a function $\widetilde{\sigma}$ locally defined in $\mathcal{N}_{\bar{U}}$ mapping games to equilibria, with $\widetilde{\sigma}(\bar{U}) = \bar{\sigma}$, is innocuous; we assume $\hat{U}$ is close enough to $\bar{U}$ so that $\hat{U}$ belongs to $\mathcal{N}_{\bar{U}}$, the domain of $\widetilde{\sigma}$. Consider the homotopy function $H\colon \mathcal{N}_{\bar{\sigma}} \times [0,1] \to \mathbb{R}^{|S|\sum_{i\in I}|A_i|}$ defined by

$$H(\sigma, \tau) = (1-\tau) f(\sigma, \bar{U}) + \tau f(\sigma, \hat{U}).$$

The existence and uniqueness of a path $\sigma\colon [0,1] \to \mathcal{N}_{\bar{\sigma}}$ satisfying $H(\sigma(\tau), \tau) = 0$, for all $\tau \in [0,1]$, is ensured by our results. Moreover, such a path yields us the desired new equilibrium since $\sigma(1) \in \mathbf{Equil}(\hat{U}) \cap \mathcal{N}_{\bar{\sigma}}$. Of course, the question is then how to find the path. An alternative is to solve the ordinary differential equation

$$\frac{d\sigma}{d\tau}(\tau) = -\Big[\frac{df}{d\sigma}\big(\sigma(\tau), (1-\tau)\bar{U} + \tau\hat{U}\big)\Big]^{-1}\big(f(\sigma(\tau), \hat{U}) - f(\sigma(\tau), \bar{U})\big)$$

with initial condition $\sigma(0) = \bar{\sigma}$. Consult, for example, Garcia and Zangwill (1981) for additional discussion.

If $\mathbf{Equil}(\hat{U})$ has several elements, one could ask whether $\hat{\sigma} = \sigma(1)$, being the only equilibrium close to the original observed equilibrium, is a sensible prediction for the new dynamic model. Implicit in the previous paragraph's exercise is the assumption that the selection procedure that lead the players to play $\bar{\sigma}$ in original game $\bar{U}$ is continuous, so that when the game is only slightly modified to $\hat{U}$, players will be lead to play the only equilibrium $\hat{\sigma}$ which is close enough to the original equilibrium $\bar{\sigma}$. Once again, here we touch a topic about which not much is known in dynamic stochastic games. To provide some preliminary and informal support to $\hat{\sigma}$ as a sensible prediction to the dynamic model $\hat{U}$, suppose that $\bar{\sigma}$ is a sink (so an asymptotically stable point) of some learning process. Suppose that the flow characterizing the learning process is a sufficiently continuous function of the game; for example, suppose that the flow is derived from best replies which, as a consequence of the the maximum theorem, are continuous functions of the parameters of the game. Since $\bar{\sigma}$ is a sink and the learning rule is continuous in the game, $\hat{\sigma}$ must be a sink to the new learning process induced by the game $\hat{U}$; $\hat{\sigma}$ therefore is an asymptotically stable point of the learning process of the new game. Moreover, as $\bar{\sigma}$, the starting point of the learning process, is sufficiently close to the asymptotically stable point $\hat{\sigma}$, $\hat{\sigma}$ is the result of the learning process starting at $\bar{\sigma}$.[7]

7 Purification of Regular Equilibria

This section presents our main purification results. We begin by introducing incomplete information into our baseline dynamic stochastic games. After presenting the purification theorem, we briefly discuss some of its implications.

7.1 Perturbed Dynamic Stochastic Games

Below we consider a slightly different version of the model studied in the previous sections. Following Harsanyi (1973a), we now assume that in every period t player i receives a shock $\eta_i^t \in \mathbb{R}^{|A|}$ before choosing his action. The shock η_i is known to player i but not to his rivals. The private shocks are independent across players and periods, and drawn from a probability distribution μ_i. We assume that μ_i is absolutely continuous with respect to Lebesgue measure on $\mathbb{R}^{|A|}$. The period payoff function of player i is

$$u_i(a, s) + \eta_i(a),$$

where $\eta_i(a)$ denotes the a component of η_i. We refer to this private information game as the *perturbed dynamic stochastic game.*

The equilibrium notion is Bayesian Markov perfect equilibrium. We, however, are not interested in equilibrium strategies but in equilibrium distributions. A profile $\sigma = (\sigma_i)_{i \in I} \in \Sigma$ is a *Markov perfect equilibrium distribution* (equilibrium distribution, for short) if there exists a function $\bar{V}_i \colon S \to \mathbb{R}$, for each $i \in I$, such that

[7] See Fudenberg and Levine (1998) for an exposition of the theory of learning in games.

3. for all i and all s

$$\bar{V}_i(s) =$$

$$\int \left(\max_{a_i \in A_i} u_i((a_i, \sigma_{-i}(\cdot, s)), s) + \eta_i(a_i, \sigma_{-i}(\cdot, s)) + \delta_i \sum_{s' \in S} \bar{V}_i(s') q(s'; (a_i, \sigma_{-i}(\cdot, s)), s) \right) \mu_i(d\eta_i),$$

and

4. for all s, the probability profile $\sigma(\cdot, s) = (\sigma_i(\cdot, s)_{i \in I}$ is *consistent*, meaning that there exists a profile $b = (b_i(\eta_i))_{i \in I}$ which is a Bayesian equilibrium of the incomplete information game where, given an action profile $a \in A$, player i's payoff is

$$u_i(a, s) + \eta_i(a) + \delta_i \sum_{s' \in S} \bar{V}_i(s') q(s'; a, s)$$

and such that for all $a_i \in A_i$

$$\sigma_i(a_i, s) = \int_{\{\eta_i | b_i(\eta_i) = a_i\}} \mu_i(d\eta_i)$$

Conditions 3 and 4 are similar in spirit to Conditions 1 and 2 introduced in Section 2.2. The main difference is that now there is incomplete information and therefore the equilibrium concept in the static game is Bayesian. The equilibrium notion to the perturbed dynamic stochastic game corresponds to a Bayesian Markov perfect equilibrium; Escobar (2007) ensures the existence of such an equilibrium in this and related models.

7.2 Purification: Convergence and Approachability

We are now ready to explore how good of an approximation to the original dynamic stochastic game the perturbed game is. More precisely, we consider, for all $i \in I$, a sequence of probability distributions of private shocks $(\mu_i^n)_{n \in \mathbb{N}}$ converging to a mass point at $0 \in \mathbb{R}^{|A|}$. We ask whether the corresponding sequence of perturbed games has equilibrium distributions that are getting closer (in a sense to be specified below) to the equilibria of the original (unperturbed) game. To simplify the exposition we define $\mathbf{Equil}^n(U)$ to be the set of equilibrium distributions of the perturbed game when players' private shocks are drawn from $\mu^n = (\mu_i^n)_{i \in I}$.

The following result, whose proof is in the Appendix, shows that as the private shocks vanish, any converging sequence of equilibrium distributions of perturbed games converges to an equilibrium of the unperturbed stochastic game.

Proposition 10 (Convergence) *Suppose that, for all $i \in I$, $(\mu_i^n)_{n \in \mathbb{N}}$ converges to a mass point at $0 \in \mathbb{R}^{|A|}$ as $n \to \infty$. Suppose further that $\sigma^n \in \mathbf{Equil}^n(U)$ converge to $\bar{\sigma}$ as $n \to \infty$. Then, $\bar{\sigma} \in \mathbf{Equil}(U)$.*

The following is the main result of this section. It shows that any equilibrium of the original (unperturbed) game can be approximated by equilibrium distributions of nearby perturbed games.

Theorem 11 (Approachability) *Suppose that, for all $i \in I$, $(\mu_i^n)_{n\in\mathbb{N}}$ converges to a mass point at $0 \in \mathbb{R}^{|A|}$ as $n \to \infty$. Let $\bar{\sigma}$ be a regular equilibrium of game U. Then for all $\bar{\epsilon} > 0$ and all large enough n, there exists $\sigma^n \in \mathbf{Equil}^n(U)$ that is within a distance $\bar{\epsilon}$ of $\bar{\sigma}$.*

It is considerably more difficult to obtain lower hemi-continuity results such as Theorem 11 than closure results such as Proposition 10. The proof of Theorem 11 relies on arguments previously presented by Govindan, Reny, and Robson (2003). The two key properties satisfied by regular equilibria that are exploited to generalize Govindan, Reny, and Robson's (2003) proof are strong stability and quasi-strictness; see Section 8.2 for details.

7.3 Discussion

These results offer a compelling interpretation to behavior strategy equilibria in dynamic games. According to this interpretation, we should think of the original complete information game as an idealization, a game that does not capture every payoff relevant consideration. Moreover, at each decision node, each player has small private information. As shown by our analysis, a behavior strategy equilibrium of the complete information dynamic game can be seen as an equilibrium distribution of the more comprehensive dynamic game of incomplete information.

Our analysis also answers the problem of formulation robustness posed by recent numerical procedures solving dynamic stochastic games. It turns out that the introduction of incomplete information into a dynamic model facilitates the computation of equilibrium strategies (Pakes and McGuire 1994, Pakes and McGuire 2001). Our results show that the choice between complete and incomplete information is largely one of convenience.

The equilibrium strategies studied in Examples 2 and 3, being regular, can be seen through the lens of our purification results. For the exit game of Example 2, Doraszelski and Satterthwaite (2007) adapt Pakes and McGuire's (1994) routines and compute the equilibrium strategies of a family of perturbed exit game as incomplete information vanishes. The analysis applied to the strategies introduced in Example 3 implies that these strategies are robust to the introduction of incomplete information.

It is also worth noting the implications of our results to a more general class of repeated games. As demonstrated by Fudenberg and Maskin (1986), the set of feasible payoffs in a repeated game is quite large. Is is then natural to ask whether introducing small incomplete information into the repeated game model provides a sensible way to select some of the strategies supporting those payoffs.

To answer this question, we proceed as in Example 3 by embedding the repeated game model into our dynamic setting. Consider an arbitrary repeated game satisfying the full dimensionality assumption of Fudenberg and Maskin (1986), an individually strictly rational payoff vector $v \in \mathbb{R}^{|I|}$, and a discount factor such that the folk theorem strategies introduced by Fudenberg and Maskin (1986) constitute an equilibrium. Recall that the folk theorem strategies call for a punishment phase after a deviation of, say, player i, where his rivals minimax player i for $T_i \in \mathbb{N}$ periods, and then a carrot phase where game players are rewarded for carrying out the punishment phase. To represent these strategies, we consider a dynamic stochastic game with $1 + \sum_{i \in I} T_i + |I|$, where the first state amounts for the cooperation phase, and for each player i, T_i states amount for i's punishment phase and 1 state amounts for the subsequent carrot phase. If public randomization is required in some phase, then it is possible to expand the state space to represent the public randomization. Moreover, after inspection of the proof of Theorem 2 in Fudenberg and Maskin (1986) it is easy to see that the incentives can be taken strict, both on- and off-path. So, folk theorem strategies, being strict, are purifiable and therefore robust to vanishing incomplete information.[8] We conclude that at least for the set of equilibrium payoffs studied, introducing vanishing incomplete information into a repeated game model cannot yield a sensible selection procedure.

8 Proofs

8.1 Proof of Theorem 4

8.1.1 The Invertibility Lemma

The (a, s) component of the vector $QV_i(\cdot, \sigma) \in \mathbb{R}^{|A||S|}$ is the future stream of discounted payoffs when the current state is s and the action profile played is a. Therefore, the (a, s) component of $U_i + \delta_i QV_i(\cdot, \sigma)$ is the expected discounted payoff of player i when the system is in state s and the strategy profile is a. Moreover, $V_i(\cdot, \sigma)$ defined by (3.3) assumes the form

$$V_i(\cdot, \sigma) = \Big(\sum_{t=0}^{\infty} (\delta_i)^t (\mathcal{P}_\sigma Q)^t \mathcal{P}_\sigma \Big) U_i.$$

Now,

$$\begin{aligned} U_i + \delta_i QV_i(\cdot, \sigma) &= \Big(\mathbb{I}_{|A||S|} + \delta_i Q \sum_{t=0}^{\infty} (\delta_i)^t (\mathcal{P}_\sigma Q)^t \mathcal{P}_\sigma \Big) U_i \\ &= \Big(\sum_{t=0}^{\infty} (\delta_i)^t (Q\mathcal{P}_\sigma)^t \Big) U_i \\ &= \Big(\mathbb{I}_{|A||S|} - \delta_i Q \mathcal{P}_\sigma \Big)^{-1} U_i. \end{aligned}$$

[8] We are implicitly assuming that on- and off-path play is in pure strategies. This entails a loss of generality because the more general equilibrium strategies allowing mixed strategy minimax phases (Theorem 5 in Fudenberg and Maskin (1986)) being nontrivially behavioral, are not strict. We have not been able to establish whether or not they are regular.

All these manipulations are justified because, by definition of Σ^{ϵ}, the matrices involved are invertible.

The previous observations are simple, but important enough to be highlighted in a lemma.

Lemma 12 (Invertibility Lemma) *For all i and all $\sigma \in \Sigma^{\epsilon}$, the matrix $\left(\mathbb{I}_{|A||S|} - \delta_i Q \mathcal{P}_{\sigma}\right) \in \mathbb{R}^{|A||S| \times |A||S|}$ is invertible. Therefore, the function $U_i \mapsto U_i + \delta_i Q V_i(\cdot, \sigma)$ that maps per period payoffs to payoffs in the reduced one-shot game is linear and invertible.*

In other words, by knowing the family of normal form games induced by a continuation profile $\sigma \in \Sigma^{\epsilon}$, we can know the payoff structure U of the dynamic stochastic game.

8.1.2 Two Useful Lemmata

We now present two intermediate lemmata. A corollary to the second lemma characterizing the dimension of the equilibrium graph is then presented and discussed.

To present these results, we need to introduce additional notation. For given product sets $B^*, C^* \in \prod_{i \in I} \prod_{s \in S} 2^{|A_i|}$, define $G(B^*, C^*)$ as the set of games having some equilibrium σ exhibiting best replies B^* and carriers C^*. Formally,

$$G(B^*, C^*) = \left\{U \mid \exists \sigma \in \mathbf{Equil}(U) \text{ such that for all } i, B_i(\sigma, \cdot) = B_i^* \text{ and } C_i(\sigma, \cdot) = C_i^*\right\},$$

where $B_i(\sigma, s)$ and $C_i(\sigma, s)$ are sets introduced in Section 6.1. We also define the set $I(B^*, C^*)$ of games having an irregular equilibrium σ for which for all i, $B_i(\sigma, \cdot) = B_i^*$ and $C_i(\sigma, \cdot) = C_i^*$.

Lemma 13 shows that for almost all games their equilibria are quasi-strict (see Section 6.1 for the definition of quasi-strict equilibrium). The proof proceeds as follows. We first derive a set of necessary conditions characterizing a game $\bar{U}$ and an equilibrium $\bar{\sigma}$ which is not quasi-strict. These equations, written as $M(\bar{\sigma}, \bar{U}) = 0$ in the proof, are linearly independent (as shown below in Claim 1) and therefore we can derive a locally defined function mapping strategies and *some* components of the payoff vector U to the whole vector U. It is then shown that the set of games close to $\bar{U}$ having some equilibrium near to $\bar{\sigma}$ (in the proof, this set is $G^{\bar{\sigma}, \bar{U}}(B^*, C^*)$) has a small dimension and is therefore negligible (Claim 2). The result follows applying this logic to each possible pair $(\bar{U}, \bar{\sigma})$ in a properly chosen way.

Lemma 13 *For all $B^* \neq C^*$, $\lambda(G(B^*, C^*)) = 0$. In particular, $\lambda(I(B^*, C^*)) = 0$.*

Proof. For all $i \in I$ and $s \in S$, consider an element $a_i^s \in C_i^*(s)$. Consider a game $\bar{U}$ having some equilibrium $\bar{\sigma}$ such that $B_i(\bar{\sigma}, \cdot) = B_i^*$ and $C_i(\bar{\sigma}, \cdot) = C_i^*$. For all $i \in I$, all $a_i \in A_i$, and all $s \in S$, $\mathcal{P}^s_{a_i, \bar{\sigma}_{-i}} \left(\mathbb{I}_{|A||S|} - \delta_i Q \mathcal{P}_{\bar{\sigma}}\right)^{-1} \bar{U}_i$ is the payoff that i obtains in state s when it plays a_i in the current period, given that its rivals follow the strategy $\bar{\sigma}_{-i}$ and i follows the strategy $\bar{\sigma}_i$ from the next period on. By assumption, for all $a_i \in B_i^*(s)$, $\mathcal{P}^s_{a_i,\bar\sigma_{-i}}\big(\mathbb{I}_{|A||S|} - \delta_i Q\mathcal{P}_{\bar\sigma}\big)^{-1}\bar U_i$ equals $\mathcal{P}^s_{a_i^s,\bar\sigma_{-i}}\big(\mathbb{I}_{|A||S|} - \delta_i Q\mathcal{P}_{\bar\sigma}\big)^{-1}\bar U_i$. In other words, $(\mathcal{P}^s_{a_i,\bar\sigma_{-i}} - \mathcal{P}^s_{a_i^s,\bar\sigma_{-i}})\big(\mathbb{I}_{|A||S|} - \delta_i Q\mathcal{P}_{\bar\sigma}\big)^{-1}\bar U_i$ is 0 when $a_i \in B_i^*(s)$. For a fixed i, consider the equations $(\mathcal{P}^s_{a_i,\bar\sigma_{-i}} - \mathcal{P}^s_{a_i^s,\bar\sigma_{-i}})\big(\mathbb{I}_{|A||S|} - \delta_i Q\mathcal{P}_{\bar\sigma}\big)^{-1}\bar U_i = 0, \;$ for all $a_i \in B_i^*(s)$. These equations can be written in matrix form as

$$\bar{\mathcal{P}}_{i,\bar\sigma}\big(\mathbb{I}_{|A||S|} - \delta_i Q\mathcal{P}_{\bar\sigma}\big)^{-1}\bar U_i = 0, \tag{8.1}$$

where $\bar{\mathcal{P}}_{i,\bar\sigma} \in \mathbb{R}^{\sum_{s\in S}(|B_i^*(s)|-1)\times|A||S|}$ is the vertical concatenation of the row vectors $\mathcal{P}^s_{a_i,\bar\sigma_{-i}} - \mathcal{P}^s_{a_i^s,\bar\sigma_{-i}}$ over all the pairs (a_i, s) such that $a_i \in B_i^*(s) \setminus \{a_i^s\}$. We consider the whole system of equations

$$M(\bar\sigma, \bar U) \equiv \begin{pmatrix} \bar{\mathcal{P}}_{1,\bar\sigma}\big(\mathbb{I}_{|A||S|} - \delta_1 Q\mathcal{P}_{\bar\sigma}\big)^{-1}\bar U_1 \\ \vdots \\ \bar{\mathcal{P}}_{|I|,\bar\sigma}\big(\mathbb{I}_{|A||S|} - \delta_{|I|} Q\mathcal{P}_{\bar\sigma}\big)^{-1}\bar U_{|I|} \end{pmatrix} = 0, \tag{8.2}$$

which imposes the whole set of indifference conditions for all players i.

Claim 1 *Consider any game $\bar U$ and any equilibrium $\bar\sigma$ such that for all i, $B_i(\bar\sigma,\cdot) = B_i^*$. Then, the derivative $\frac{\partial M(\bar\sigma,\bar U)}{\partial \bar U}$ has rank $\sum_{i\in I}\sum_{s\in S}\big(|B_i^*(s)| - 1\big)$.*

This claim, whose proof can be found in the Appendix, follows by exploiting the diagonal structure of $\frac{\partial M(\bar\sigma,\bar U)}{\partial \bar U}$ and by noting that each $\bar{\mathcal{P}}_{i,\bar\sigma}$ has full rank and, as a consequence of Lemma 12, so does $\bar{\mathcal{P}}_{i,\bar\sigma}\big(\mathbb{I}_{|A||S|} - \delta_i Q\mathcal{P}_{\bar\sigma}\big)^{-1}$.

Consider any pair $(\bar U, \bar\sigma)$ as in the above claim. As a consequence of the implicit function theorem, there exist open sets $\mathcal{N}^1 \subseteq \mathbb{R}^{|I||A||S| - \sum_{i\in I}\sum_{s\in S}(|B_i^*(s)|-1)}$, $\mathcal{N}^2 \subseteq \mathbb{R}^{|I||A||S|}$ and $\mathcal{N} \subseteq \Sigma^\epsilon$, and a function $\Phi\colon \mathcal{N}^1 \times \mathcal{N} \to \mathcal{N}^2$, where $\bar U \in \mathcal{N}^1 \times \mathcal{N}^2$ (properly ordered) and $\bar\sigma \in \mathcal{N}$, such that for all $(U^1, \sigma) \in \mathcal{N}^1 \times \mathcal{N}$, $U^2 = \Phi(U^1, N^1)$ is the only solution in $\mathcal{N}^2$ to $M(\sigma, U^1, U^2) = 0$. We define the function $H(U^1, \sigma) = (U^1, \Phi(U^1, \sigma))$. In order to highlight the dependence of these objects on $\bar\sigma, \bar U$, we write $\mathcal{N}^1_{\bar\sigma,\bar U}$, $\mathcal{N}^2_{\bar\sigma,\bar U}$, $\mathcal{N}_{\bar\sigma,\bar U}$, and $H_{\bar\sigma,\bar U}$, respectively. We also assume, without loss of generality, that $\mathcal{N}^1_{\bar\sigma,\bar U}$, $\mathcal{N}^2_{\bar\sigma,\bar U}$ and $\mathcal{N}_{\bar\sigma,\bar U}$ are balls having rational centers and radii. Define the sets

$$G^{\bar\sigma,\bar U}(B^*, C^*) = \\ \Big\{U \in \mathcal{N}^1_{\bar\sigma,\bar U} \times \mathcal{N}^2_{\bar\sigma,\bar U} \mid \sigma \in \mathbf{Equil}(U) \cap \mathcal{N}_{\bar\sigma,\bar U} \text{ with } B_i(\sigma,\cdot) = B_i^* \text{ and } \sigma \in A(C^*)\Big\},$$

where $A(C^*) = \{\sigma \in \Sigma \mid C_i(\sigma,\cdot) = C_i^* \text{ for all } i\}$, and

$$P^{\bar\sigma,\bar U}(B^*, C^*) = \\ \Big\{U \in \mathcal{N}^1_{\bar U,\bar\sigma} \times \mathcal{N}^2_{\bar U,\bar\sigma} \mid \exists (U^1, \sigma) \in \mathcal{N}^1_{\bar U,\bar\sigma} \times \big(A(C^*) \cap \mathcal{N}_{\bar U,\bar\sigma}\big) \text{ such that } H_{\bar U,\bar\sigma}(U^1,\sigma) = U\Big\}.$$

Note that $G^{\bar\sigma,\bar U}(B^*, C^*) \subseteq P^{\bar\sigma,\bar U}(B^*, C^*)$. This logic allows us to construct the sets $G^{\bar\sigma,\bar U}(B^*, C^*)$ and $P^{\bar\sigma,\bar U}(B^*, C^*)$ for each $(\bar\sigma, \bar U)$ such that $\bar\sigma$ is an equilibrium for game $\bar U$ having $C(\bar\sigma) = C^*$ and

$B(\bar{\sigma}) = B^*$.

Claim 2 *Consider any game $\bar{U}$ and an equilibrium $\bar{\sigma}$ of $\bar{U}$ such that $C(\bar{\sigma}) = C^*$ and $B(\bar{\sigma}) = B^*$. Then $\lambda\big(P^{\bar{\sigma},\bar{U}}(B^*, C^*)\big) = 0$.*

The proof of this claim is as follows. Note first that $\dim(\mathcal{N}^1_{\bar{\sigma},\bar{U}}) = |I||A||S| - \sum_{i\in I}\sum_{s\in S}(|B_i^*(s)| - 1)$, $\dim(A(C^*)\cap\mathcal{N}_{\bar{\sigma},\bar{U}}) = \sum_{i\in I}\sum_{s\in S}\big(|C_i^*(s)|-1\big)$. Therefore $\dim(\mathcal{N}^1_{\bar{\sigma},\bar{U}}\times\big(A(C^*)\cap\mathcal{N}_{\bar{\sigma},\bar{U}})\big) = |I||A||S| - \sum_{i\in I}\sum_{s\in S}|B_i^*(s)| + \sum_{i\in I}\sum_{s\in S}|C_i^*(s)| < |I||A||S|$ when $B_i^*(s) \neq C_i^*(s)$ for some i and some s. Since $P^{\bar{\sigma},\bar{U}}(B^*, C^*) = H_{\bar{\sigma},\bar{U}}\big(\mathcal{N}^1_{\bar{\sigma},\bar{U}} \times \big(A(C^*)\cap\mathcal{N}_{\bar{\sigma},\bar{U}}\big)\big)$, the result follows.

By construction, we can find a countable sequence of sets $P^{n'}$, $n' \in \mathbb{N}$, such that $G(B^*, C^*) \subseteq \cup_{n'\in\mathbb{N}} P^{n'}$ with $\lambda(P^{n'}) = 0$. The result follows by noting that the union of measure zero sets is measure zero as well ■

The proof of Lemma 13 resembles proofs given for normal form games by Harsanyi (1973a) and van Damme (1991). The main difference is that we are not able to define M globally. We indeed think that such M cannot be globally defined in our dynamic game setting. We therefore analyze the system of equations locally, and apply this construction to a countable set of games and equilibria.

The previous lemma allows us to dispose of all games in which some best reply is not being played. Now, we introduce a result which permits us to analyze games where all best replies are played in equilibrium. Before we turn to the next result, we introduce the set of completely mixed distributions in Σ^ϵ as

$$\widetilde{\Sigma} = \Big\{\sigma \in \Sigma^\epsilon \mid \sigma_i(a_i, s) > 0 \text{ for all } i, a_i, s\Big\}.$$

The lemma below shows that, when restricted to $\widetilde{\Sigma}$, the Jacobian of f with respect to the pair (σ, U) has full rank. Its proof, which is similar to that of Claim 1, exploits Lemma 12 and the diagonal structure of the Jacobian.

Lemma 14 *Consider a strategy profile $\sigma \in \widetilde{\Sigma}$. Then $\frac{\partial f(\sigma,U)}{\partial(\sigma,U)}$ has full rank $|S|\sum_{i\in I}|A_i|$.*

Proof. For $s \in S$ and $a_i \in A_i \setminus \{a_i^s\}$, define the row (a_i, s) of $P_i^*(\sigma)$ by

$$\sigma_i(a_i, s)\big(\mathcal{P}^s_{a_i,\sigma_{-i}} - \mathcal{P}^s_{a_i^s,\sigma_{-i}}\big).$$

Note that $P_i^*(\sigma) \in \mathbb{R}^{(|A_i|-1)|S|\times|A||S|}$. Therefore, the components of f associated to player i can be represented by

$$f_i(\sigma, U) = \begin{pmatrix} \sum_{a_i\in A_i}\sigma_i(a_i, s_1) - 1 \\ \vdots \\ \sum_{a_i\in A_i}\sigma_i(a_i, s_K) - 1 \\ P_i^*(\sigma)(\mathbb{I}_{|A||S|} - \delta_i Q\mathcal{P}_\sigma)^{-1}U_i \end{pmatrix} \in \mathbb{R}^{|A_i||S|}$$

The derivative of the first $|S|$ components of f_i with respect to σ_i takes the form

$$\left(\begin{array}{ccc|ccc|cc|ccc}
 & \sigma_i(\cdot,s_1) & & & \sigma_i(\cdot,s_2) & & & \dots & & \sigma_i(\cdot,s_K) & \\
\hline
1 & \dots & 1 & 0 & \dots & 0 & & & 0 & \dots & 0 \\
0 & \dots & 0 & 1 & \dots & 1 & 0 & \dots & & \dots & 0 \\
\vdots & & & & & & & & & & \vdots \\
0 & \dots & & & & & & & 1 & \dots & 1
\end{array}\right)$$

This matrix, denoted X_i, has rank $|S|$. The derivative of the first $|S|$ components of f_i with respect to (σ_{-i}, U) is 0.

Consider now the components of f_i associated to $a_i \neq a_i^s$. The derivative of those components with respect to U takes the form

$$\left(\begin{array}{c|c|c|c|c|c|c|c}
U_1 & U_2 & \dots & U_{i-1} & U_i & U_{i+1} & \dots & U_{|I|} \\
\hline
0 & 0 & \dots & 0 & P_i^*(\sigma)(\mathbb{I}_{|A||S|} - \delta_i Q \mathcal{P}_\sigma)^{-1} & 0 & \dots & 0
\end{array}\right)$$

But $P_i^*(\sigma)$ has full rank $|S|(|A_i| - 1)$. Indeed, each row (a_i, s), with $a_i \neq a_i^s$, of $P_i^*(\sigma)$ is zero in all the components (a', s') such that either $s' \neq s$ or a' does not contain a_i. Moreover, for each row a_i, s of $P_i^*(\sigma)$ there must be some nonzero component (a', s), with a' containing a_i. This proves that $P_i^*(\sigma)$ has full rank and therefore so does $P_i^*(\sigma)(\mathbb{I}_{|A||S|} - \delta_i Q \mathcal{P}_\sigma)^{-1}$.

These calculations show that, up to permutation of rows,

$$\frac{\partial f(\sigma, U)}{\partial(\sigma, U)} = \left(\begin{array}{c|c|c|c|c|c|c|c}
\sigma_1 & \sigma_2 & \dots & \sigma_{|I|} & U_1 & U_2 & \dots & U_{|I|} \\
\hline
X_1 & 0 & 0 & \dots & 0 & 0 & \dots & 0 \\
0 & X_2 & 0 & \dots & 0 & 0 & \dots & 0 \\
\vdots & & \ddots & & & & \vdots & \\
0 & \dots & 0 & X_{|I|} & 0 & \dots & \dots & 0 \\
\hline
 & & & & Z_1 & 0 & \dots & 0 \\
 & & & & 0 & Z_2 & \dots & 0 \\
Y_1 & Y_2 & & Y_{|I|} & \vdots & \vdots & \ddots & \vdots \\
 & & & & 0 & & \dots & Z_{|I|}
\end{array}\right),$$

where $Z_i = P_i^*(\sigma)(\mathbb{I}_{|A||S|} - \delta_i Q \mathcal{P}_\sigma)^{-1}$ and $Y_i \in \mathbb{R}^{|S| \sum_{i \in I}(|A_i|-1) \times |A_i||S|}$, for all i. This permits us to deduce that the rank of $\frac{\partial f(\sigma,U)}{\partial(\sigma,U)}$ is $|S| \sum_{i \in I} |A_i|$. ■

Lemma 14 implies the following version of the structure theorem.

Corollary 15 *The equilibrium graph has the same dimension as the space of games.*

This corollary generalizes the observation made by Govindan and Wilson (2001) for normal form games to stochastic games. If we were interested only in finiteness results, this lemma and the transversality theorem would yield those results almost immediately. (To obtain finiteness results, it is enough to restrict attention to completely mixed equilibria; therefore, there is no need to establish Lemma 13.) Our analysis therefore provides a much easier finiteness proof than that of Haller and Lagunoff (2000). Our arguments also provide a finiteness proof for normal form games that cannot be deemed as more complicated than the simple finiteness proof given by Govindan and Wilson (2001), who exploit the theory of semi-algebraic sets.[9]

8.1.3 Proof of Theorem 4

We employ the following result from differential topology known as the transversality theorem.

Theorem 16 *Let $\mathcal{O} \subseteq \mathbb{R}^q$ be open and $L\colon \mathcal{O} \times \mathbb{R}^s \to \mathbb{R}^n$ be C^k where $k \geq 1 + \max\{0, q-n\}$. Assume that the Jacobian of L has rank n for all $(x,y) \in \mathcal{O} \times \mathbb{R}^s$ such that $L(x,y) = 0$. Now fix $y = \bar{y}$. Then $L(\cdot, \bar{y}) : \mathcal{O} \to \mathbb{R}^n$ is regular for almost all $\bar{y}$. In other words, for almost all $\bar{y}$, the Jacobian $\frac{\partial L(x,\bar{y})}{\partial x}$ has rank n for all $x \in \mathcal{O}$ such that $L(x,\bar{y}) = 0$.*

The transversality theorem is a generalization of the well known Sard's theorem. See Abraham and Robbin (1967) and Guillemin and Pollack (1974) for additional results and technical details.

We are now in a position to prove Theorem 4. Denote by $\bar{I}$ the set of all games having some irregular equilibrium. Then

$$\bar{I} = \bigcup_{C^* \subseteq B^*} I(B^*, C^*).$$

Since there exists a finite number of sets B^*, C^* such that $C^* \subseteq B^*$, it is enough to show that for all such sets $\lambda(I(B^*, C^*)) = 0$. If $B^* \neq C^*$, this follows from Lemma 13.

If $B^* = C^*$, then an implication of equation (6.1) in Proposition 6 is that $\partial f / \partial \sigma$ has full rank if and only if so does the matrix obtained by crossing out the rows and columns associated to components (a_i, s) such that $a_i \notin B_i^*(s)$. This submatrix is itself a Jacobian matrix associated to the optimality conditions imposed on a game having a completely mixed equilibrium. We can therefore assume without loss of generality that $B^*(s) = C^*(s) = A$ for all s. Then,

$$I(B^*, C^*) \subseteq \left\{ U \in \mathbb{R}^{|I||A||S|} \mid \exists \sigma \in \widetilde{\Sigma} \text{ such that } f(\sigma, U) = 0 \text{ and } \frac{\partial f(\sigma, U)}{\partial \sigma} \text{ is singular} \right\}.$$

From Lemma 14, $\frac{\partial f(\sigma,U)}{\partial(\sigma,U)}$ has full rank for all $(\sigma, U) \in \widetilde{\Sigma} \times \mathbb{R}^{|I||A||S|}$. The transversality theorem therefore implies that for almost all U, $f(\cdot, U)|_{\widetilde{\Sigma}} \colon \widetilde{\Sigma} \to \mathbb{R}^{|S| \sum_{i \in I} |A_i|}$ is regular.

[9]However, note that if we were to extend Govindan and Wilson's (2001) tools to stochastic games, we would need to additionally prove the semi-algebraicity of the map f.

8.2 Proof of Theorem 11

8.2.1 Continuation Values

Given Condition 4 in the definition of equilibrium distribution, Condition 3 can be replaced by

$$\bar{V}_i(s) = u_i(\sigma(\cdot,s),s) + \delta_i \sum_{s' \in S} \bar{V}_i(s')q(s';\sigma(\cdot,s),s) + \int \sum_{a_i \in A_i} \eta_i(a_i,\sigma_{-i}(\cdot,s))\mathbb{1}_{\{\eta_i | a_i \text{ is optimal given } s\}}\mu_i(d\eta_i) \tag{8.3}$$

where the event $\{\eta_i \mid a_i \text{ is optimal given } s\}$ simply summarizes the maximization problem inside the integral in Condition 3. It is not immediate to see that the last term in equation (8.3) can be written solely as a function of exogenous variables and the equilibrium distribution σ, without reference to equilibrium strategies; consult Aguirregabiria and Mira (2007) for details. Now, equation (8.3), as equation (3.1), is a well defined expression for $\sigma \in \Sigma$ which may not be an equilibrium distribution. Moreover, for all $\sigma \in \Sigma$, equation (8.3), seen as a system of $|S|$ equations in the unknowns $\bar{V}_i(\cdot) \in \mathbb{R}^{|S|}$, has a unique solution, which we denote $\bar{V}_i(\cdot,\sigma) \in \mathbb{R}^{|S|}$. The analysis presented by Aguirregabiria and Mira (2007) implies that the vector $\bar{V}_i(\cdot,\sigma)$ can be written as

$$\bar{V}_i(\cdot,\sigma) = \big(\mathbb{I}_{|S|} - \delta_i \mathcal{P}_\sigma Q\big)^{-1}\big(\mathcal{P}_\sigma U_i + e_i(\sigma)\big), \tag{8.4}$$

where $e_i(\sigma)$ properly accounts for the term $\int \sum_{a_i \in A_i} \eta_i(a_i,\sigma_{-i}(\cdot,s))\mathbb{1}_{\{\eta_i | a_i \text{ is optimal given } s\}}\mu_i(d\eta_i)$ in equation (8.3). Being the integral of a bounded function, $e_i(\sigma)$ is a continuous function of $\sigma \in \Sigma$ and therefore so is $\bar{V}_i(\cdot,\sigma)$.

For all i, all $s \in S$ and all $\sigma \in \Sigma$, define $g_i(\cdot,s,\sigma) \in \mathbb{R}^{|A_i|}$ as the unique distribution consistent with player i's maximization given the continuation function $\bar{V}_i(\cdot,\sigma)$. More formally, given σ, consider the unique solution $b_i^\sigma \colon S \times \mathbb{R}^A \to A_i$ of the problem

$$\max_{a_i \in A_i} \; u_i(a_i,\sigma_{-i}(\cdot,s),s) + \eta_i(a_i,\sigma_{-i}(\cdot,s),s) + \delta_i \sum_{s' \in S} \bar{V}_i(s',\sigma)q(s';a_i,\sigma_{-i}(\cdot,s),s).$$

In other words, $b_i^\sigma(s,\eta_i)$ is the optimal decision of player i when the state is s, its private shock is η_i, given that his rivals are following the strategy profile σ_{-i} and given that from tomorrow on σ_i is consistent with i's future actions. Define $g_i(a_i,s,\sigma) = \int_{\{\eta_i | b_i^\sigma(s,\eta_i) = a_i\}} \mu_i(\eta_i)$, which is a continuous function of σ.

8.2.2 Proof of Theorem 11

We employ the following result from algebraic topology.

Proposition 17 (Govindan, Reny, and Robson (2003)) *Suppose that O is a bounded, open set in $\mathbb{R}^{\bar{m}}$ and $h, f \colon \bar{O} \to \mathbb{R}^{\bar{m}}$ are continuous, where $\bar{O}$ denotes the closure of O. Further, suppose that f is continuously differentiable on O, that x_0 is the only zero of f in O and that the Jacobian* *of f at x_0 has full rank. If, for every $t \in [0,1]$, the function $th + (1-t)f$ has no zero on the boundary of O, then h has a zero in O.*

Denote by $\bar{V}_i^n$ and g_i^n, the corresponding continuation and consistent distribution functions defined for the dynamic stochastic game with payoff perturbations $(\mu_i^n)_i$. For each $n \in \mathbb{N}$, define the function $h^n \colon \mathbb{R}^{|S|\sum_{i\in I}|A_i|} \to \mathbb{R}^{|S|\sum_{i\in I}|A_i|}$ by

$$h^n_{i,a_i,s}(\sigma) = \begin{cases} \sum_{a_i \in A_i} \sigma_i(a_i,s) - 1 & \text{if} \quad a_i = a_i^s, \\ g^n_{i,a_i}(s;\sigma) - \sigma_i(a_i,s) & \text{if} \quad a_i \neq a_i^s. \end{cases}$$

Just for concreteness, $\|\cdot\|$ denotes the Euclidean norm. Since $\bar{\sigma}$ is regular, the argument presented in the proof of Proposition 7 shows the existence of an open set O such that the following conditions, referred to hereafter as C1-C5, are satisfied:

C1 $\bar{\sigma} \in O$

C2 $\|\bar{\sigma} - \sigma\| < \bar{\epsilon}$, $\sigma \in O$

C3 $\bar{\sigma}$ is the only zero of $f(\cdot, U)$ in O

C4 For all i, s, a_i, if $\bar{\sigma}_i(a_i,s) > 0$, then for any $\sigma \in O$, $\sigma_i(a_i,s) > 0$

C5 For all i, s, a_i, if $\mathcal{U}_i^U(a_i,s,\bar{\sigma}) - \mathcal{U}_i^U(a_i^s,s,\bar{\sigma}) < 0$, then for any $\sigma \in O$, $\mathcal{U}_i^U(a_i,s,\sigma) - \mathcal{U}_i^U(a_i^s,s,\sigma) < 0$

To prove Theorem 11 it is enough to find a zero of h^n in O for n large enough. As a consequence of Proposition 17, the following result yields the desired result.

Lemma 18 *For all large enough n, and all $t \in [0,1]$, the function $th^n + (1-t)f(\cdot,U)$ has no zero on the boundary of O.*

Proof. Suppose not, that is, consider a sequence $(t^n)_{n\in\mathbb{N}}$ converging to $\bar{t}$ in $[0,1]$ and a sequence $(\sigma^n)_{n\in\mathbb{N}}$, contained in the boundary of O, converging to $\hat{\sigma}$, such that σ^n is a zero of $t^n h^n + (1-t)f(\cdot,U)$ for all n. We present two preliminary claims.

Claim 3 *Consider i, a_i and s such that $\mathcal{U}_i^U(a_i,s,\hat{\sigma}) - \mathcal{U}_i^U(a_i^s,s,\hat{\sigma}) < 0$. Then, $\hat{\sigma}_i(a_i,s) = 0$.*

The proof of this claim is as follows. Equations (3.3) and (8.4) imply that $\bar{V}_i^n(\sigma^n) \to V_i(\hat{\sigma})$, and so $g_i^n(a_i,s,\sigma^n)$ goes to 0. Therefore,

$$\begin{aligned} 0 =& \lim_n t^n h^n_{i,s,a_i}(\sigma^n) + (1-t^n) f_{i,s,a_i}(\sigma^n, U) \\ =& \hat{t}(-\hat{\sigma}_i(a_i,s)) + (1-\hat{t}) f_{i,s,a_i}(\hat{\sigma}, U) \\ =& -\hat{\sigma}_i(a_i,s)\Big(\hat{t} + (1-\hat{t})(\mathcal{U}_i^U(a_i^s,s,\hat{\sigma}) - \mathcal{U}_i^U(a_i,s,\hat{\sigma}))\Big). \end{aligned}$$

It follows that $\hat{\sigma}_i(a_i, s) = 0$.

Claim 4 *For all i and s, $\hat{\sigma}(\cdot, s)$ is a probability distribution and there exist i, s, and $a_i \neq a_i^s$ such that $f_{i,a_i,s}(\hat{\sigma}, U) = 0$.*

To see why this result holds, note that if $\hat{\sigma}_i(a_i, s) < 0$, C4 implies that $\bar{\sigma}_i(a_i, s) = 0$. Since $\bar{\sigma}$ is quasi-strict, $\mathcal{U}_i^U(a_i, s, \bar{\sigma}) - \mathcal{U}_i^U(a_i^s, s, \bar{\sigma}) < 0$ and, from C5, $\mathcal{U}_i^U(a_i, s, \hat{\sigma}) - \mathcal{U}_i^U(a_i^s, s, \hat{\sigma}) < 0$. Claim 3 shows that $\hat{\sigma}_i(a_i, s) = 0$, which constitutes a contradiction. Therefore, $\hat{\sigma}_i(a_i, s) \geq 0$ for all i, a_i, s. Further, because $h^n_{i,a_i^s,s} \equiv f_{i,a_i^s,s}$, we deduce that $f_{i,a_i^s,s}(\sigma^n, U) = 0$ for all n, so $\hat{\sigma}_i(\cdot, s)$ is a well defined probability distribution on actions. From C3, $\hat{\sigma}$ cannot be a zero of $f(\cdot, U)$. So there must exist $i, a_i \neq a_i^s, s$ such that $f_{i,a_i,s}(\hat{\sigma}, U) \neq 0$.

With these two preliminary claims at hand, we now complete the proof. Fix i, a_i, s as in Claim 4 for the rest of the proof.

First note that a_i^s cannot belong to $B_i(\hat{\sigma}, s)$. Indeed, if it did, then $\mathcal{U}_i^U(a_i, s, \hat{\sigma}) - \mathcal{U}_i^U(a_i^s, s, \hat{\sigma}) \leq 0$ and, from Claim 3, $f_{i,a_i,s}(\hat{\sigma}, U) = 0$. Since $\bar{\sigma}$ is quasi-strict and as a consequence of C4, $\hat{\sigma}_i(a_i', s) > 0$ for all $a_i' \in B_i(\bar{\sigma}, s)$. C5 implies that $B_i(\hat{\sigma}, s) \subseteq B_i(\bar{\sigma}, s)$ and therefore $\hat{\sigma}_i(a_i', s) > 0$ for all $a_i' \in B_i(\hat{\sigma}, s)$. Consequently,

$$\sum_{a_i' \in B_i(\hat{\sigma}, s)} f_{i,a_i',s}(\sigma^n, U) > 0. \tag{8.5}$$

For $a_i' \notin B_i(\hat{\sigma}, s)$, by definition $g_i^n(a_i', s, \sigma^n) \to 0$, so that

$$\sum_{a_i' \in B_i(\hat{\sigma}, s)} g_i^n(a_i', s, \sigma^n) \to 1.$$

Because $\hat{\sigma}_i(a_i^s, s) > 0$, and $\hat{\sigma}(\cdot, s)$ is a probability distribution, $\sum_{a_i' \in B_i(\hat{\sigma}, s)} \hat{\sigma}_i(a_i', s) < 1$. Therefore

$$\sum_{a_i' \in B_i(\hat{\sigma}, s)} h^n_{i,a_i',s}(\sigma^n) > 0, \tag{8.6}$$

for n large enough. But equations (8.5) and (8.6) imply that $t^n h^n + (1 - t^n) f(\cdot, U)$ is not zero at σ^n, which is a contradiction. This completes the proof. ■

9 Appendix: Omitted Proofs

This Appendix contains the proofs of some supporting results.

Proof of Claim 1. For each i, the matrix $\bar{\mathcal{P}}_{i,\bar{\sigma}}$ has full rank $\sum_{s \in S}(|B_i^*(s)| - 1)$. To see this, note that for each (a_i, s) where $a_i \in B_i^*(s) \backslash \{a_i^s\}$, the row vector $\mathcal{P}^s_{a_i,\bar{\sigma}_{-i}} - \mathcal{P}^s_{a_i^s,\bar{\sigma}_{-i}} \in \mathbb{R}^{|A| \times |S|}$ contains a zero in all those components (a', s') where either $s' \neq s$ or a_i is not contained in a' (by this we mean that there is no a_{-i} such that $(a_i, a_{-i}) = a'$). This row vector also contains some nonzero term in some component (a', s) where a' contains a_i; indeed, $\sum_{a' \text{ contains } a_i} \bar{\sigma}_{-i}(a' \setminus a_i, s) = \sum_{a_{-i}} \bar{\sigma}_{-i}(a_{-i}, s) = 1$. This proves the desired rank result.

The derivative of M with respect to $\bar{U}$ takes the diagonal form

$$\frac{\partial M(\bar{\sigma}, \bar{U})}{\partial \bar{U}} = \begin{pmatrix} \bar{\mathcal{P}}_{1,\bar{\sigma}}\left(\mathbb{I}_{|A||S|} - \delta_1 Q \mathcal{P}_{\bar{\sigma}}\right)^{-1} & 0 & \dots & 0 \\ 0 & \bar{\mathcal{P}}_{2,\bar{\sigma}}\left(\mathbb{I}_{|A||S|} - \delta_2 Q \mathcal{P}_{\bar{\sigma}}\right)^{-1} & \vdots & 0 \\ \vdots & & & \\ 0 & 0 & \dots & \bar{\mathcal{P}}_{|I|,\bar{\sigma}}\left(\mathbb{I}_{|A||S|} - \delta_{|I|} Q \mathcal{P}_{\bar{\sigma}}\right)^{-1} \end{pmatrix}$$

As already shown, $\bar{\mathcal{P}}_{i,\bar{\sigma}}$ has rank $\sum_{s\in S}(|B_i^*(s)| - 1)$ while $\left(\mathbb{I}_{|A||S|} - \delta_i Q \mathcal{P}_{\bar{\sigma}}\right)$ is a squared invertible matrix. These imply that the diagonal matrix above has rank $\sum_{i\in I}\sum_{s\in S}\left(|B_i^*(s)| - 1\right)$. ■

Proof of Proposition 10. For n big enough, for all $s \in S$ and all $i \in I$, $C_i(\bar{\sigma}, s) \subseteq C_i(\sigma^n, s)$. Denote by $\bar{V}_i^n(\cdot, \sigma^n)$ the related continuation value functions, and note that equations (3.3) and (8.4) imply that $\bar{V}_i^n(\cdot, \sigma^n) \to V_i(\cdot, \sigma)$. By definition, any $a_i \in C_i(\sigma^n, s)$ maximizes the expression

$$u_i(a_i, \sigma_{-i}^n(\cdot, s), s) + \delta_i \sum_{s' \in S} \bar{V}_i^n(s') q(ds'; a_i, \sigma_{-i}^n(\cdot, s), s) + \eta_i(a_i, \sigma_{-i}^n(\cdot, s))$$

for all $\eta_i \in R^n(i, s) \subseteq \mathbb{R}^{|A|}$, where $\mu_i^n(R^n(i, s)) > 0$. Fixing $a_i \in C_i(\bar{\sigma}, s)$ and taking $n \to \infty$, it follows that $a_i \in B_i(\bar{\sigma}, s)$. In other words, $\bar{\sigma} \in \mathbf{Equil}(U)$. ■

References

Abraham, R., and J. Robbin (1967): *Transversal mappings and flows.* W. A. Benjamin, New York.

Acemoglu, D., and J. Robinson (2001): "A theory of political transition," *American Economic Review*, 91(4), 938–963.

Aguirregabiria, V., and P. Mira (2007): "Sequential estimation of dynamic discrete games," *Econometrica*, 75(1), 1–54.

Bajari, P., L. Benkard, and J. Levin (2007): "Estimating dynamic models of imperfect competition," *Econometrica*, 75(5), 1331–1370.

Bergemann, D., and J. Välimäki (1996): "Learning and strategic pricing," *Econometrica*, 64(5), 1125–1149.

Bernheim, D., and D. Ray (1989): "Markov perfect equilibria in altruistic growth economies with production uncertainty," *Journal of Economic Theory*, 47(1), 195–202.

Bhaskar, V. (2000): "The robustness of repeated game equilibria to incomplete payoff information," Working paper, University of Essex, Colchester.

Bhaskar, V., G. Mailath, and S. Morris (2007): "Purification in the infinitely-repeated prisoners' dilemma, second version," Working paper 07-024, Penn Institute for Economic Research, Philadelphia.

Doraszelski, U., and A. Pakes (2007): "A framework for applied dynamic analysis in IO," in *Handbook of Industrial Organization*, ed. by M. Armstrong, and R. Porter, vol. 3, pp. 1887–1966. North-Holland, Amsterdam.

Doraszelski, U., and M. Satterthwaite (2007): "Computable Markov-perfect industry dynamics: Existence, purification, and multiplicity," Working paper, Harvard University, Cambridge.

Erdem, E., and J. Tybout (2003): "Trade policy and industrial sector reponses: Using evolutionary models to interpret the evidence," Working paper no. 9947, NBER, Cambridge.

Ericson, R., and A. Pakes (1995): "Markov-perfect industry dynamics: A framework for empirical work," *Review of Economic Studies*, 62, 53–82.

Escobar, J. (2007): "Existence of pure and behavior strategy stationary Markov equilibrium in dynamic stochastic games," Working paper, Stanford University, Stanford.

Federgruen, A. (1978): "On N-person stochastic games with denumerable state space," *Advances in Applied Probability*, 10, 452–471.

Fink, A. (1964): "Equilibrium points of stochastic noncooperative games," *Journal of Science Hiroshima University Series A*, 28, 89–93.

Fudenberg, D., and D. Levine (1998): *The theory of learning in games*. MIT Press, Cambridge.

Fudenberg, D., and E. Maskin (1986): "The folk theorem in repeated games with discounting or with incomplete information," *Econometrica*, 54(3), 533–554.

Garcia, C., and W. Zangwill (1981): *Pathways to solutions, fixed points, and equilibria*. Prentice-Hall.

Goettler, R., C. Parlour, and U. Rajan (2005): "Equilibrium in a dynamic limit order market," *Journal of Finance*, 60(5), 2149–2192.

Govindan, S., P. Reny, and A. Robson (2003): "A short proof of Harsanyi's purification theorem," *Games and Economic Behavior*, 45(2), 369–374.

Govindan, S., and R. Wilson (2001): "Direct proofs of generic finiteness of Nash equilibrium outcomes," *Econometrica*, 69(3), 765–769.

Guillemin, V., and A. Pollack (1974): *Differential topology*. Prentice Hall, Englewood Cliffs.

Haller, H., and R. Lagunoff (2000): "Genericity and Markovian behavior in stochastic games," *Econometrica*, 68(5), 1231–1248.

Harsanyi, J. (1973a): "Games with randomly disturbed payoffs: A new rationale for mixed-strategy equilibrium points," *International Journal of Game Theory*, 2(1), 1–23.

——— (1973b): "Oddness of the number of equilibrium points: A new proof," *International Journal of Game Theory*, 2(4), 235–250.

Herings, J., and R. Peeters (2004): "Stationary equilibria in stochastic games: Structure, selection, and computation," *Journal of Economic Theory*, 118, 32–60.

KADYRZHANOVA, D. (2006): "The leader-bias hypothesis: Corporate control dynamics in industry equilibrium," Working paper, Columbia University, New York.

KOJIMA, M., A. OKADA, AND S. SHINDOH (1985): "Strongly stable equilibrium points of n-person noncooperative games," *Mathematics of Operations Reserach*, 10, 650–663.

MASKIN, E., AND J. TIROLE (2001): "Markov perfect equilibrium, I: Observable actions," *Journal of Economic Theory*, 100, 191–219.

PAKES, A., AND P. MCGUIRE (1994): "Computing Markov-perfect Nash equilibria: Numerical implications of a dynamic differentiated product model," *Rand Journal of Economics*, 25(4), 555–589.

——— (2001): "Stochastic algorithms, symmetric Markov perfect equilibrium, and the "curse" of dimensionality," *Econometrica*, 69(5), 1261–1281.

PAKES, A., M. OSTROVSKY, AND S. BERRY (2004): "Simple estimators for the parameters of discrete dynamic games (with entry/exit examples)," Working paper no. 10506, NBER, Cambridge.

——— (2006): "Simple estimators for the parameters of discrete dynamic games (with entry/exit examples)," *Rand Journal of Economics*, forthcoming.

PESENDORFER, M., AND P. SCHMIDT-DENGLER (2003): "Identification and estimation of dynamic games," Working paper no. 9726, NBER, Cambridge.

SHAPLEY, L. (1953): "Stochastic games," *Proceedings of the National Academy of Sciences*, 39, 1095–1100.

SOBEL, M. (1971): "Noncooperative stochastic games," *The Annals of Mathematical Statistics*, 42(6), 1930–1935.

STOKEY, N., AND E. LUCAS, R. WITH PRESCOTT (1989): *Recursive methods in economic dynamics*. Harvard University Press, Cambridge.

VAN DAMME, E. (1991): *Stability and perfection of Nash equilibria*. Springer, Heidelberg, 2 edn.

WHITT, W. (1980): "Representation and approximation of noncooperative sequential games," *SIAM Journal of Control and Optimization*, 18(1), 33–48.

WU, W., AND J. JIANG (1962): "Essential equilibrium points of n-person non-cooperative games," *Scientia Sinica*, 11(10), 1307–1322.