



HIER

Harvard Institute of Economic Research

Discussion Paper Number 2156

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

by

Ulrich Doraszelski
and
Juan F. Escobar

April 2008

Harvard University
Cambridge, Massachusetts

This paper can be downloaded without charge from:
<http://www.economics.harvard.edu/journals/hier2008>

The Social Science Research Network Electronic Paper Collection:
<http://ssrn.com/abstract=1120819>

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

Ulrich Doraszelski
Harvard University[†]

Juan F. Escobar
Stanford University[‡]

April 14, 2008

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, strong stability, essentiality, purifiability, estimation, computation, repeated games.

*We are greatly indebted to Michaela Draganska, Drew Fudenberg, Dino Gerardi, Hari Govindan, Paul Grieco, Matt Jackson, Ken Judd, Carlos Lever, Jon Levin, Michael Ostrovsky, Martine Quinzii, Ilya Segal, Andy Skrzypacz, and Bob Wilson for comments and suggestions. Doraszelski gratefully acknowledges the hospitality of the Hoover Institution during the academic year 2006/07 and financial support from the National Science Foundation under Grant No. 0615615.

[†]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

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 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.

While several results in the literature guarantee the 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 a dynamic environment. This paucity of knowledge 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, and raises a number of crucial concerns for both theoretical and applied work on dynamic stochastic games.

First, it is not known whether slight changes in the parameters of a dynamic stochastic game cause slight changes in its Markov perfect equilibria or severely alter the nature of the interactions among players. This concern is especially pertinent if the researcher uses modern econometric techniques such as Aguirregabiria and Mira (2007), Bajari, Benkard, and Levin (2007), Judd and Su (2008), Pakes, Ostrovsky, and Berry (2006), and Pesendorfer and Schmidt-Dengler (2008) to structurally estimate the underlying primitives of the game and therefore recovers the payoffs with estimation error. It then becomes critical that slight changes in the payoffs do not completely reshape equilibrium behavior.

A second and related concern is that it is not clear that comparative statics are well defined in dynamic environments, in particular when it comes to taking derivatives of the endogenous variables with respect to an exogenous variable. Applied work also often aims to conduct counterfactual experiments or policy simulations (e.g., Benkard 2004, Collard-Wexler 2005, Dube, Hitsch, and Manchanda 2005, Gowrisankaran and Town 1997, Ryan 2005). These exercises present an especially troubling problem when a dynamic stochastic game has multiple equilibria because in this case little is known about which equilibrium will be played after a change to the system has occurred. What happens depends on how players adjust to the change.

A third concern relates to the estimation of primitives and the computation of equilibria. 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 a game of incomplete

information 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, making them the natural starting point for structural estimation. Yet, to date it is not known whether the choice between complete and incomplete information is largely one of convenience or whether the strategic interactions among players are sensitive to the different formulations of a dynamic economic problem.

A fourth and final concern is the interpretation of behavior strategy equilibria in dynamic environments. In a mixed-strategy equilibrium, at each decision node, a player is indifferent among several pure actions. It is then natural to ask what compels the player to randomize precisely as mandated by the equilibrium. For normal form games Harsanyi's (1973a) celebrated purification theorem provides an elegant answer to this question, namely that a mixed-strategy equilibrium can be seen as a pure-strategy equilibrium of a nearby game of incomplete information. In general classes of dynamic economic models, in contrast, it is not known whether purifiability is the norm or the exception.

The goal of this paper is to help settle these concerns 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 showing that regularity is a generic property of Markov perfect equilibria. More formally, we identify a dynamic stochastic game with its period payoffs and show that the set of games having Markov perfect equilibria that all are regular has full Lebesgue measure. While regularity is a purely mathematical concept, it paves the way to a number of economically meaningful properties. A consequence of the fact that all Markov perfect equilibria of almost all dynamic stochastic games are regular is that almost all games have a finite number of Markov perfect equilibria that are locally isolated. Moreover, these equilibria are essential and strongly stable and are therefore robust to slight changes in payoffs. Finally, they all admit purification and can therefore be obtained as limits of equilibria of dynamic stochastic games of incomplete information as random payoff fluctuations become vanishingly small. In sum, this paper shows how to extend several of the most fundamental 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) from static to dynamic environments.

Our main insight is that, holding fixed the value of future play, the strategic situation that the players face in a given state of the dynamic system is akin to a normal form game. Consequently, a Markov perfect equilibrium of a dynamic stochastic game must satisfy the conditions for a Nash equilibrium of a certain reduced one-shot game. We exploit these conditions to derive a system of equations, $f(\sigma) = 0$, that must be satisfied by any Markov perfect equilibrium σ . We say that the equilibrium is regular if the Jacobian of f with respect to σ , $\frac{\partial f(\sigma)}{\partial \sigma}$, 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 advanced 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 the main reason that we are able to obtain our genericity, stability, and purification results. Indeed, because we view a dynamic stochastic game as a family of interrelated (and endogenous) normal form games, we are able to “import” many of the techniques that have been used to prove these results in the context of normal form games.

The proof of our main genericity result builds on the seminal work of Harsanyi (1973b). The presence of nontrivial dynamics, however, introduces nonlinearities that preclude us from simply applying Harsanyi’s (1973b) construction. Two insights are the key to our proof. First, the map that relates a dynamic stochastic game to the payoff matrices of the family of induced normal form games underlying our regularity notion is linear and invertible. These properties are evident if players fully discount the future and the dynamic stochastic game collapses to a set of disjoint normal form games, but they are less than obvious in the presence of nontrivial dynamics. Second, in order to study the regularity of f , we do not have to work with an alternative globally defined function, relate its regular points to those of f , and then apply Sard’s theorem to the alternative function. In a departure from the standard treatment in the literature on normal form games (Harsanyi 1973b, 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—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 dynamic stochastic games, our focus on one-shot deviations has the additional advantage that it permits us to generalize several other major results for normal form games besides generic finiteness to dynamic stochastic games.¹

We demonstrate that regular Markov perfect equilibria are robust to slight changes 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 on normal form games to our dynamic setting, namely essentiality and strong stability. 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 changes in payoffs. We show that regular equilibria are strongly stable and, therefore, essential. This result in combination with our main genericity result yields the

¹As 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 can be deduced by elaborating on our arguments.

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.

These stability properties ensure that slight changes in the parameters of a dynamic stochastic game do not severely alter the nature of the interactions among players. In addition, they lay the foundations for comparative statics. Because the map from payoffs to equilibria is differentiable, differentiable comparative statics are well defined, at least for a small change to the system. We also offer some guidance for the particularly difficult situation when a dynamic stochastic game has multiple equilibria. Under a variety of learning processes, our main stability result allows us to single out the equilibrium that is likely to be played after a small change to the system has occurred. We finally discuss how to compute this equilibrium using so-called homotopy or path-following methods.

Next we show that regular Markov perfect equilibria admit purification, thereby extending Harsanyi's (1973a) celebrated purification theorem from normal form games to dynamic stochastic games. We perturb a dynamic stochastic game by assuming that, at each decision node, a player's payoffs are subject to random fluctuations that are known to the player but not to his rivals. We demonstrate that any regular Markov perfect equilibrium of the original complete information game can be obtained as the limit of equilibria of the perturbed game of incomplete information as payoff fluctuations become vanishingly small. Hence, one can view the original game of complete information as an idealization—a limit—of nearby games with a small amount of payoff uncertainty. The proof of our main purification result generalizes arguments previously presented by Govindan, Reny, and Robson (2003) in the context of normal form games. That we are able to do so once again shows the power of our regularity notion.

Our main purification result suggests that the choice between complete and incomplete information in formulating a dynamic economic problem is largely one of convenience, at least in situations where the random payoff fluctuations are deemed small. It also provides a convincing interpretation of behavior strategy equilibria in dynamic stochastic games because in the approximating equilibrium a player is no longer indifferent among several pure actions but instead has a strictly optimal pure action for almost all realizations of his payoffs.

We finally advance the study of purifiability in dynamic environments. While the purifiability of equilibrium behavior in dynamic environments has received some attention recently, the literature so far has studied but a small number of particular examples. Bhaskar (1998, 2000) and Bhaskar, Mailath, and Morris (2007) provide examples of nonpurifiable equilibria in which strategies depend on payoff irrelevant variables. In contrast, our main purification result shows that equilibria in which strategies depend only on the payoff relevant history are generically purifiable. We discuss in more detail how our results relate to their examples later in the paper. For now we just note that, to the best of our knowledge, this paper is the first to establish the purifiability of equilibrium

²Maskin 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.

behavior in a general class of dynamic economic models.

In sum, in this paper we develop 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 equilibria. These equilibria are essential and strongly stable. Moreover, they all admit purification. The key to obtaining these results is our notion of regularity which is based on the insight that, holding fixed the value of future play, the strategic situation that the players face in a given state of the dynamic system is akin to a normal form game. By viewing a dynamic stochastic game as a family of induced normal form games, we are able to make a rich body of literature on normal form games useful for the analysis of dynamic environments.

The remainder of this paper is organized as follows. Section 2 sets out the model and equilibrium concept. Section 3 introduces our notion of regularity and illustrates it with two examples. Section 4 states the main genericity result and discusses its implications for the finiteness of the equilibrium set. Section 5 presents stability properties and Section 6 our main purification result. Section 7 contains the proofs of our main genericity and purification results. Some supporting arguments have been relegated to the Appendix.

2 Model

In this section we set up the model and define our notion of equilibrium. We further describe the total payoff that a player receives in a dynamic stochastic game.

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 expected 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.³ Time is discrete and the horizon is infinite.⁴

The game proceeds as follows. The dynamic system starts at time $t = 0$ 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 over 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

³The case of state-dependent action sets can be analyzed using exactly the same arguments that we develop below (see also Example 3). 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 were not surprised if suitable modifications of our arguments apply to these models as well.

⁴Our results also apply to finite-horizon dynamic stochastic games.

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 $u_i: A \times S \rightarrow \mathbb{R}$ is the period payoff function of player i . 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. In the next round at time $t = 1$, 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.

A dynamic stochastic game is characterized by a tuple

$$\langle S, (A_i, u_i, \delta_i)_{i \in I}, q, \bar{q} \rangle,$$

where $\delta_i \in [0, 1[$ is the discount factor of player i that is used to compute his total payoff as the expected net present value of his period payoffs. In the remainder of this paper, 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.

2.2 Markov Perfect Equilibria

Roughly speaking, a Markov perfect equilibrium is a subgame perfect equilibrium in which the strategies depend only on the payoff relevant history. Below we provide a precise definition of our equilibrium concept and an alternative characterization that is key to the subsequent 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 $\Sigma_i = \Delta(A_i)^{|S|}$ 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 .

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 **Equil**(U). The nonemptiness of **Equil**(U) has long been established in the literature (see, e.g., Fink 1964).

We next provide an alternative characterization of equilibrium that exploits the recursive structure of the model. A strategy profile $\sigma = (\sigma_i)_{i \in I}$ is a Markov perfect equilibrium if and only if (i) for all $i \in I$ there exists a function $V_i: S \rightarrow \mathbb{R}$ such that for all $s \in S$

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

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

$$u_i(a, s) + \delta_i \sum_{s' \in S} V_i(s') q(s'; a, s). \quad (2.2)$$

The function $V_i: S \rightarrow \mathbb{R}$ in equation (2.1) is the equilibrium value function for player i . $V_i(s)$ is the expected net present 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 .

Our alternative characterization of equilibrium 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. Consequently, an equilibrium of the dynamic stochastic must induce a Nash equilibrium in a certain reduced one-shot game. The payoff to player i in this game as given in equation (2.2) is the sum of his period payoff and his appropriately discounted continuation value. Note that, given continuation values, equation (2.2) can be used to construct the entire payoff matrix of the normal form game that players face in state s .

While simple, the observation that a dynamic stochastic game can be studied by analyzing a family of normal form games is the key to the subsequent analysis. It suggests to define a notion of regularity with reference to the induced normal form games. We formalize this idea in Section 3. The obvious difficulty that we have to confront is that the induced normal form games are endogenous in that the payoffs depend on the equilibrium of the dynamic stochastic game.

2.3 Notation and Continuation Values

Before defining our notion of regularity, we introduce some notation and further describe the total payoff that a player receives in a dynamic stochastic game.

In what follows we consider not only equilibria but also deviations from equilibrium strategies. We thus have to know the value of continued play given an arbitrary strategy profile $\sigma \in \Sigma$. In fact, since Σ is not an open set, we work mostly with the set Σ^ϵ . We construct Σ^ϵ to be open in $\mathbb{R}^{\sum_{i \in I} |A_i| |S|}$ and to strictly contain Σ . The construction of Σ^ϵ is detailed in the Appendix. Here we just note that Σ^ϵ has elements that are not strategies.

To facilitate the subsequent analysis we introduce some notation. Enumerate action profiles $a^1, \dots, a^{|A|}$ such that

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

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

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

We further define the matrix $\mathcal{P}_\sigma \in \mathbb{R}^{|S| \times |A||S|}$ as

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

In the remainder of the paper, we write \mathcal{P}_σ^s to denote row s of \mathcal{P}_σ . Consider the delta Dirac strategy $\mathbb{1}_{a_i}$ that puts probability 1 on action $a_i \in A_i$ in all states $s \in S$. For all $\sigma_{-i} \in \Sigma_{-i}^\epsilon$, $\mathcal{P}_{(\mathbb{1}_{a_i}, \sigma_{-i})}$ is well defined and we simply denote it by $\mathcal{P}_{a_i, \sigma_{-i}}$. Finally, \mathbb{I}_r denotes the $\mathbb{R}^{r \times r}$ identity matrix.

Using this notation, the value of continued play given an arbitrary profile $\sigma \in \Sigma^\epsilon$ is

$$V_i^U(\cdot, \sigma) = \left(\sum_{t=0}^{\infty} (\delta_i)^t (\mathcal{P}_\sigma Q)^t \mathcal{P}_\sigma \right) U_i = \left(\mathbb{I}_{|S|} - \delta_i \mathcal{P}_\sigma Q \right)^{-1} \mathcal{P}_\sigma U_i, \quad (2.3)$$

where the inversion is justified by the construction of Σ^ϵ . We interpret $V_i^U(s, \sigma)$ as the expected net present value of the stream of payoffs to player i if the dynamic system is currently in state s and play is according to σ . Our notation emphasizes that these continuation values also depend on the payoff vector U .

3 Regularity

In this section we first define our notion of regularity. Then we illustrate our definition with two examples.

3.1 Regular Markov Perfect Equilibria

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 with the payoffs in equation (2.2). Since an equilibrium of the dynamic stochastic must induce a Nash equilibrium of this game, one-shot deviations cannot be profitable.

To make this idea precise, define the function $\mathcal{U}_i : A_i \times S \times \Sigma^\epsilon \rightarrow \mathbb{R}$, where

$$\mathcal{U}_i(a_i, s, \sigma) = u_i((a_i, \sigma_{-i}(\cdot, s)), s) + \delta_i \sum_{s' \in S} V_i^U(s', \sigma) q(s'; (a_i, \sigma_{-i}(\cdot, s)), s). \quad (3.1)$$

$\mathcal{U}_i(a_i, s, \sigma)$ is the expected net present value of the stream of payoffs to player i if the current state is s , his rivals play according to σ_{-i} , and player i chooses action a_i in the current period and then plays according to σ_i from the subsequent period on. If $\sigma \in \Sigma$ is an equilibrium, then one-shot deviations cannot be profitable. Formally, if $\sigma \in \Sigma$ is an equilibrium, then, for all $i \in I$ and $s \in S$,

$$\sigma_i(a_i, s) > 0 \text{ implies } a_i \in \arg \max_{a'_i \in A_i} \mathcal{U}_i(a'_i, s, \sigma). \quad (3.2)$$

These are the conditions for a Nash equilibrium in the normal form game in state s as induced by the value of continued play $V_i^U(\cdot, \sigma)$.

Consider a collection of actions $a_i^s \in A_i$ for all $i \in I$ and $s \in S$. We think of a_i^s as a reference action for player i in state s . Define the (i, a_i, s) component of the function $f : \Sigma^\epsilon \times \mathbb{R}^{|I||A||S|} \rightarrow \mathbb{R}^{|S| \sum_{i \in I} |A_i|}$ by

$$f_{i, a_i, s}(\sigma, U) = \begin{cases} \sum_{a_i \in A_i} \sigma_i(a_i, s) - 1 & \text{if } a_i = a_i^s, \\ \sigma_i(a_i, s) \left(\mathcal{U}_i(a_i, s, \sigma) - \mathcal{U}_i(a_i^s, s, \sigma) \right) & \text{if } a_i \neq a_i^s. \end{cases} \quad (3.3)$$

Condition (3.2) 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

$$f(\sigma, U) = 0. \quad (3.4)$$

Note that equation (3.4) is necessary but not sufficient for an equilibrium. Further, since \mathcal{P}_σ is continuously differentiable as a function of $\sigma \in \Sigma^\epsilon$, so is $V_i^U(\cdot, \sigma)$. Therefore, f is continuously differentiable as a function of $\sigma \in \Sigma^\epsilon$. For future reference, we note that f is also continuously differentiable as a function of $(\sigma, U) \in \Sigma^\epsilon \times \mathbb{R}^{|I||A||S|}$.

We are now ready to define our notion of regularity.

Definition 2 *A Markov perfect equilibrium σ of a dynamic stochastic game U 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 $a_i^s \in A_i$ such that $\sigma(a_i^s, s) > 0$ for all $i \in I$ and $s \in S$. An equilibrium is irregular if it is not regular.*

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 reduced one-shot games. This observation, while simple, permits us to make a rich body of literature on normal form games useful for the analysis of dynamic environments.

3.2 Examples

Below we provide two examples of dynamic stochastic games that have regular equilibria.

Example 3 (Exit Game) Our first example is a simple version of the model of industry dynamics proposed by Ericson and Pakes (1995) (see also Doraszelski and Satterthwaite 2007). Consider an industry with two firms. The set of players is $I = \{1, 2\}$ and the set of states is $S = \{(1, 1), (1, 0), (0, 1), (0, 0)\}$, where the state $s = (s_1, s_2)$ indicates whether firm i is in the market ($s_i = 1$) or out of the market ($s_i = 0$).

While we assumed in Section 2.1 that action spaces are state invariant, we use this example to illustrate how easy it is to consider the more general setting. In particular, to simplify the exposition, we assume that the only nontrivial decision a firm has to make is whether or not to exit if the market is a duopoly in state $(1, 1)$; in all other states a firm has no choice but to preserve the status quo. More formally, firm i 's action set is $\{\text{exit}, \text{stay}\}$ in state $s = (1, 1)$ but a singleton in all other states. Thus, all states other than state $(1, 1)$ are absorbing.

If the market is a duopoly in state $(1, 1)$, then each firm receives a period payoff π^D . If the market is a monopoly in state $(1, 0)$ or $(0, 1)$, then the monopolist receives a period payoff π^M and its competitor receives nothing. Neither firm receives anything in state $(0, 0)$. Finally, if a firm is in the market but decides to exit it, then it receives a scrap value ϕ regardless of what the other firm does. We assume that

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

where $\delta \in]0, 1[$ is the common discount factor. Hence, while a monopoly is viable, a duopoly is not. The duopolists are thus caught up in a war of attrition.

The exit game has three equilibria, two of them in pure strategies. The only symmetric equilibrium, denoted by $\bar{\sigma}$, is fully characterized by the probability of exiting if the market is a duopoly:

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

Our goal here is to show that the symmetric equilibrium $\bar{\sigma}$ is regular.

To compute the value of continued play given an arbitrary strategy profile σ , we exploit the simple structure of the exit game rather than rely on equation (2.3). Since all states other than

state $(1, 1)$ are absorbing, we have $V_i(s, \sigma) = 0$ if $s_i = 0$ and $V_i(s, \sigma) = \frac{\pi^M}{1-\delta}$ if $s_i = 1$ and $s_{-i} = 0$. $V_i((1, 1), \sigma)$ is defined recursively as the unique solution to

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

The (i, a_i, s) component of f is

$$f_{i, a_i, s}(\sigma) = \begin{cases} \sigma_i(a_i, s) - 1 & \text{if } s_i \neq (1, 1), \\ \sigma_i(\text{exit}, (1, 1)) + \sigma_i(\text{stay}, (1, 1)) - 1 & \text{if } s_i \neq (1, 1) \text{ and } a_i = \text{stay}, \\ \sigma_i(\text{exit}, (1, 1)) \left\{ \pi^D + \phi - \left(\pi^D + \delta \sigma_{-i}(\text{exit}, (1, 1)) \frac{\pi^M}{1-\delta} \right. \right. \\ \left. \left. + \delta \sigma_{-i}(\text{stay}, (1, 1)) V_i((1, 1), \sigma) \right\} & \text{if } s_i \neq (1, 1) \text{ and } a_i = \text{exit}. \end{cases}$$

Computing the Jacobian of f with respect to σ and evaluating its determinant at $\bar{\sigma}$, it can be verified that the determinant is nonzero under our assumptions on the parameters, so that the symmetric equilibrium $\bar{\sigma}$ is regular. This is easiest to see if we normalize $\pi^D = 0$ to reduce the determinant to

$$-\frac{\delta^2 \phi^2 (\pi^M - (1-\delta)\phi)^4}{\left(\delta (\pi^M)^2 - (1-\delta)^2 \phi^2 \right)^2} < 0.$$

Example 4 (Repeated Prisoners' Dilemma) Our second example shows how a repeated prisoners' dilemma and its trigger strategies can be embedded into our dynamic stochastic game framework. Consider the repeated prisoners' dilemma with two players $I = \{1, 2\}$, 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

where $l, g > 0$. Let \mathcal{H} be the set of all histories of the repeated game. Define the trigger strategy $\bar{a}_i: \mathcal{H} \rightarrow A_i$ by

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

If $\delta > 1/(1+g)$, then the trigger strategy constitutes a subgame perfect equilibrium of the repeated prisoners' dilemma. Moreover, at each decision node, a player strictly prefers to play the actions prescribed by the trigger strategy.

To embed the trigger strategy equilibrium of the repeated game within our framework, define

the state space by $S = \{\text{on}, \text{off}\}$ and 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 $\bar{q}(\text{on}) = 1$. Consider the stationary Markov behavior strategy

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

It is easy to check that this strategy constitutes a Markov perfect equilibrium of the dynamic stochastic game if $\delta > 1/(1 + g)$. Moreover, at each decision node, a player strictly prefers to conform to the stationary Markov behavior strategy (we call these equilibria strict, see Section 5.1 for a formal definition). Proposition 8 therefore implies that the Markov equilibrium is regular. We return to repeated games later in the paper.

4 Genericity of Regular Equilibria

Before demonstrating that our notion of regularity is useful for characterizing the equilibrium set and the properties of its members, we show that regularity is a property that is satisfied by all equilibria of a large set of models. Our regularity notion is thus widely applicable.

Recall that we identify a dynamic stochastic game with its period payoff functions $(u_i)_{i \in I}$. 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.

Theorem 5 *For almost all games $U \in \mathbb{R}^{|I||A||S|}$, all equilibria are regular.*

The proof of Theorem 5 is detailed in Section 7.1. It 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 5.1 for a formal definition). Within this class we restrict attention to completely mixed equilibria. For these equilibria, we show that the Jacobian 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

⁵See Haller and Lagunoff (2000) for an insightful discussion of alternative genericity notions.

desired result, namely that the Jacobian 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) and further illustrated by van Damme (1991). One difference is, of course, the technical difficulties that we have to surmount when working with dynamic stochastic games instead of normal form games. The presence of nontrivial dynamics introduces nonlinearities that preclude us from simply applying Harsanyi's (1973b) construction. Indeed, the family of induced normal form games is endogenous to the equilibrium of the dynamic stochastic game.

Two insights facilitate our analysis. The first observation 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 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 7.1.2 for details).

The second observation facilitating our analysis is that, given a strategy profile σ to be followed from next period on, the map that relates a dynamic stochastic game to the payoff matrices of the family of induced normal form games is linear and invertible. To see this, consider the normal form game induced in state s . Given the action profile a , player i obtains a payoff

$$u_i(a, s) + \delta_i \sum_{s' \in S} V_i^U(s', \sigma) q(s'; a, s), \quad (4.1)$$

where we have replaced the equilibrium continuation value $V_i(\cdot)$ in equation (2.2) with $V_i^U(\cdot, \sigma)$, the value of continued play given the arbitrary strategy profile σ to be followed from next period on. The payoff to player i in equation (4.1) can be written in matrix notation as the (a, s) component of the vector $U_i + \delta_i Q V_i^U(\cdot, \sigma) \in \mathbb{R}^{|A||S|}$. Using equation (2.3), we obtain

$$\begin{aligned} U_i + \delta_i Q V_i^U(\cdot, \sigma) &= \left(\mathbb{I}_{|A||S|} + \delta_i Q \sum_{t=0}^{\infty} (\delta_i)^t (\mathcal{P}_\sigma Q)^t \mathcal{P}_\sigma \right) U_i \\ &= \left(\mathbb{I}_{|A||S|} + \delta_i Q \mathcal{P}_\sigma + \delta_i Q \delta_i (\mathcal{P}_\sigma Q) \mathcal{P}_\sigma + \delta_i Q \delta_i^2 (\mathcal{P}_\sigma Q)^2 \mathcal{P}_\sigma + \dots \right) U_i \\ &= \left(\sum_{t=0}^{\infty} (\delta_i)^t (Q \mathcal{P}_\sigma)^t \right) U_i \\ &= \left(\mathbb{I}_{|A||S|} - \delta_i Q \mathcal{P}_\sigma \right)^{-1} U_i, \end{aligned}$$

where the inversion is justified by the construction of Σ^ϵ . The following lemma summarizes the discussion so far.

Lemma 6 (Invertibility Lemma) *For all i and $\sigma \in \Sigma^\epsilon$, the matrix $(\mathbb{I}_{|A||S|} - \delta_i Q \mathcal{P}_\sigma)^{-1}$ has full rank $|A||S|$ and the map $U_i \mapsto U_i + \delta_i Q V_i^U(\cdot, \sigma)$ is linear and invertible.*

Linearity and invertibility are evident for normal form games where the term $\delta_i QV_i^U(\cdot, \sigma)$ vanishes.

In our dynamic setting, the matrix $(\mathbb{I}_{|A||S|} - \delta_i Q\mathcal{P}_\sigma)^{-1}$ is a part of the Jacobian of f with respect to U . The significance of Lemma 6 is thus that it enables us determine the rank of the Jacobian of f with respect to the pair (σ, U) , a key step in applying the transversality theorem (see Section 7.1.1 for details).

To provide a first glimpse at the power of our regularity notion, we note that 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 that consists of isolated points; therefore the equilibrium set has to be finite. We summarize in the following corollary.

Corollary 7 (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 a 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 dynamic stochastic games, 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. We discuss this point in more detail after Corollary 18. Third, and perhaps most important, our regularity notion permits us to generalize several other important results for normal form games to dynamic stochastic games.⁶ The rest of the paper illustrates this point.

5 Stability Properties of Regular Equilibria

In this section we explore the notions of strongly stable and of essential equilibria in dynamic stochastic games. Before studying these desirable 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 7. At the end of the section, we discuss the implications of our results for applied work.

⁶The regularity notion in Haller and Lagunoff (2000) does not imply quasi-strictness, a concept we formally introduce in Section 5.1. Quasi-strictness is a useful intermediate definition and establishing its connection to our regularity notion is a crucial step in deriving our main stability and purification results.

5.1 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

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

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 proposition, as well as its proof, resembles a well known result for normal form games (see, e.g., Corollary 2.5.3 in van Damme 1991).

Proposition 8 *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

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

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

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

It follows that

$$|\det(J(\sigma))| = |\det(\bar{J}(\sigma))| \prod_{i \in I} \prod_{s \in S} \prod_{a_i \notin C_i(\sigma, s)} [\mathcal{U}_i(a_i, s, \sigma) - \mathcal{U}_i(a_i^s, s, \sigma)]. \quad (5.1)$$

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, $\det(\bar{J}(\sigma)) = 1$ and $\mathcal{U}_i(a_i, s, \sigma) - \mathcal{U}_i(a_i^s, s, \sigma) < 0$ for all pairs (a_i, s) with $a_i \notin C_i(\sigma, s)$. It follows that $\det(J(\sigma)) \neq 0$ so that σ is regular. On the other hand, if the equilibrium is regular, then each of the terms on the right hand side of equation (5.1) is nonzero. Hence, $\mathcal{U}_i(a_i, s, \sigma) < \mathcal{U}_i(a_i^s, s, \sigma)$ for all pairs (a_i, s) with $a_i \notin C_i(\sigma, s)$; this corresponds to the definition of quasi-strictness. ■

5.2 Strongly Stable and Essential 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**: $\mathcal{N}_{\bar{U}} \rightarrow \mathcal{N}_{\bar{\sigma}}$ defined by **equil**(U) = **Equil**(U) \cap $\mathcal{N}_{\bar{\sigma}}$ is single-valued and continuous. In words, an equilibrium is strongly stable if the equilibrium correspondence is locally a continuous function. This definition generalizes that introduced for normal form games by Kojima, Okada, and Shindoh (1985).

Proposition 9 *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 $\tilde{\sigma}: \mathcal{N}_{\bar{U}} \rightarrow \mathcal{N}_{\bar{\sigma}}$ such that for all $U \in \mathcal{N}_{\bar{U}}$, $\tilde{\sigma}(U)$ is the unique solution $\sigma \in \mathcal{N}_{\bar{\sigma}}$ to $f(\sigma, U) = 0$. 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 $a_i \in A_i$ the following properties hold: (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 $\mathcal{U}_i(a_i, s, \bar{\sigma}) - \mathcal{U}_i(a_i^s, s, \bar{\sigma}) < 0$, then $\mathcal{U}_i(a_i, s, \sigma) - \mathcal{U}_i(a_i^s, s, \sigma) < 0$ for all $(\sigma, U) \in \mathcal{N}_{\bar{\sigma}} \times \mathcal{N}_{\bar{U}}$.

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

$$\sigma_i^U(a_i, s) > 0 \text{ for all } a_i \in C_i(\bar{\sigma}, s). \quad (5.2)$$

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

$$\mathcal{U}_i(a_i, s, \sigma^U) \leq \mathcal{U}_i(a_i^s, s, \sigma^U) \text{ with strict inequality if and only if } a_i \notin C_i(\bar{\sigma}, s). \quad (5.3)$$

Conditions (5.2) and (5.3) 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 have to satisfy $f(\sigma, U) = 0$; consequently, **equil**(U) = **Equil**(U) \cap $\mathcal{N}_{\bar{\sigma}}$ = σ^U . Since σ^U is differentiable, **equil**(U) = σ^U is differentiable (and therefore continuous). ■

For future reference we note that the proof of Proposition 9 develops an argument we exploit in the proof of our main purification result (see Section 7.2.2 for details). The idea behind the proof is to show that close enough to a regular equilibrium $\bar{\sigma}$ of \bar{U} , the system of equations $f(\sigma, U) = 0$ fully characterizes the equilibrium map **Equil**. This is not self evident for, as we noted in Section 3.1, $f(\sigma, U) = 0$ is necessary but not sufficient for $\sigma \in \mathbf{Equil}(U)$. To establish the proposition, we further invoke the quasi-strictness of a regular equilibrium, mimicking the logic of Theorem 2.5.5 in van Damme (1991).

Importantly, the proof of Proposition 9 shows that the equilibrium correspondence is locally not only a continuous but also a differentiable function.

Corollary 10 *Let $\bar{\sigma} \in \mathbf{Equil}(\bar{U})$ be regular. Then there exist open neighborhoods $\mathcal{N}_{\bar{U}}$ of \bar{U} and $\mathcal{N}_{\bar{\sigma}}$ of $\bar{\sigma}$ such that the map $\mathbf{equil}: \mathcal{N}_{\bar{U}} \rightarrow \mathcal{N}_{\bar{\sigma}}$ defined by $\mathbf{equil}(U) = \mathbf{Equil}(U) \cap \mathcal{N}_{\bar{\sigma}}$ is a differentiable function.*

To illustrate, note that the symmetric equilibrium in Example 3 is regular and therefore differentiable in the parameters of the model.

Turning to 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 proposition is immediate.

Proposition 11 *Every strongly stable equilibrium is essential.*

Theorem 5, Proposition 9, and Proposition 11 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 12 *For almost all games $U \in \mathbb{R}^{|I||A||S|}$, all equilibria are essential.*

5.3 Application: Comparative Statics and Multiple Equilibria

The literature on normal form games has forcefully argued that an equilibrium should be stable against slight changes to 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), Judd and Su (2008), Pakes, Ostrovsky, and Berry (2006), and Pesendorfer and Schmidt-Dengler (2008) to estimate the underlying primitives of a dynamic stochastic game and therefore recovers the payoffs with estimation error. It then becomes critical that slight changes in the payoffs do not severely 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 these econometric tools.

We also lay the foundations for differentiable comparative statics. Because regular equilibria are strongly stable and, as shown by Corollary 10, the locally defined equilibrium map \mathbf{equil} is differentiable, differentiable comparative statics are well defined, at least for small changes to the system.⁷

⁷Aguirregabiria and Ho (2007) conduct counterfactual experiments in a dynamic stochastic game by assuming the existence of a locally defined differentiable function that relates parameters to equilibria. While their model does not fit exactly into our framework, their analysis illustrates how the existence and differentiability of \mathbf{equil} can be exploited in applications.

While the problem of computing equilibria is beyond the scope of this paper, we briefly discuss the implications of our results for the numerical implementation of the above comparative statics exercises. Consider a game \bar{U} and an equilibrium $\bar{\sigma}$. In applied work, the fundamentals of the game \bar{U} and the equilibrium $\bar{\sigma}$ may be estimated from data. We are interested in how the equilibrium changes when we slightly change the fundamentals from \bar{U} to \hat{U} . Our results ensure the existence of a locally defined differentiable function $\mathbf{equil} : \mathcal{N}_{\bar{U}} \rightarrow \mathcal{N}_{\bar{\sigma}}$ that maps games to equilibria with $\mathbf{equil}(\bar{U}) = \bar{\sigma}$; we assume, without loss of generality, that $\mathcal{N}_{\bar{U}}$ is convex. Moreover, if \hat{U} is close enough to \bar{U} , then \hat{U} belongs to $\mathcal{N}_{\bar{U}}$, the domain of \mathbf{equil} . Consider the homotopy function $H : \mathcal{N}_{\bar{\sigma}} \times [0, 1] \rightarrow \mathbb{R}^{|\sum_{i \in I} |A_i|}$ defined by

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

Our results ensure the existence and uniqueness of a path $\sigma : [0, 1] \rightarrow \mathcal{N}_{\bar{\sigma}}$ satisfying $H(\sigma(\tau), \tau) = 0$ for all $\tau \in [0, 1]$. This path satisfies $\sigma(\tau) = \mathbf{equil}((1 - \tau)\bar{U} + \tau\hat{U})$. It is parameterized by the homotopy parameter τ and connects the “old” equilibrium $\sigma(0) = \bar{\sigma}$ at \bar{U} to a “new” equilibrium $\sigma(1) = \hat{\sigma}$ at \hat{U} . One way to compute this path is to numerically solve the ordinary differential equation

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

with initial condition $\sigma(0) = \bar{\sigma}$. See Zangwill and Garcia (1981) for a more detailed discussion of homotopy methods and Besanko, Doraszelski, Kryukov, and Satterthwaite (2007) and Borkovsky, Doraszelski, and Kryukov (2008) for an application to dynamic stochastic games.

A particular vexing problem arises when the game has multiple equilibria. In this case little is known about which equilibrium is likely be played after a change to the system has occurred. What happens depends on how players adjust to the change. See Pakes (2008) for a more detailed discussion of the multiplicity problem in applied work and Fudenberg and Levine (1998) for an exposition of the theory of learning in games. Our results offer some guidance here; in particular, under a variety of learning processes, our results suggest to single out $\sigma(1) = \hat{\sigma}$ as the equilibrium that is likely to be played after the original game \bar{U} has been slightly changed.

To provide an informal argument, consider the learning process that has given rise to the original equilibrium $\bar{\sigma}$ and note that $\bar{\sigma}$ must have been an asymptotically stable point of this learning process in the original game \bar{U} . We further assume $\bar{\sigma}$ is a sink. Suppose that the flow that characterizes the learning process is sufficiently continuous in the fundamentals of the game. While the theory of learning has not yet been developed for dynamic stochastic games, at least in the context of normal form games this is the case for a variety of learning processes, e.g., those derived from best replies. But if $\bar{\sigma}$ is a sink and the learning rule is sufficiently continuous, then $\hat{\sigma}$ must be a sink and therefore an asymptotically stable point of the learning process of the modified game \hat{U} . Moreover, as the starting point of the learning process $\bar{\sigma}$ is close enough to $\hat{\sigma}$, the learning process in the modified game will give rise to $\hat{\sigma}$.

6 Purification of Regular Equilibria

In this section we present our main purification result. We begin by introducing incomplete information into our baseline model of dynamic stochastic games. After presenting our main purification result, we briefly discuss some of its implications for theoretical and applied work.

6.1 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 of player i is

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

where $\eta_i(a)$ denotes the a component of η_i . We extend $\eta_i(\cdot)$ in the obvious way to allow for randomization over A . We refer to the private information game as the *perturbed dynamic stochastic game*; it is characterized by a tuple

$$\langle S, (A_i, u_i, \delta_i, \mu_i)_{i \in I}, q, \bar{q} \rangle.$$

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

$$\bar{V}_i(s) = \int \left(\max_{a_i \in A_i} u_i((a_i, \sigma_{-i}(\cdot, s)), s) + \eta_i(a_i, \sigma_{-i}(\cdot, s)) + \delta_i \sum_{s' \in S} \bar{V}_i(s') q(s'; (a_i, \sigma_{-i}(\cdot, s)), s) \right) d\mu_i(\eta_i) \quad (6.1)$$

and (ii) for all $s \in S$ the distribution profile $\sigma(\cdot, s) = (\sigma_i(\cdot, s))_{i \in I}$ is *consistent* with a (pure strategy) Bayesian Nash equilibrium of the incomplete information game in which player i chooses an action $a_i \in A_i$ and, given the action profile $a = (a_i)_{i \in I}$, obtains a payoff

$$u_i(a, s) + \eta_i(a) + \delta_i \sum_{s' \in S} \bar{V}_i(s') q(s'; a, s), \quad (6.2)$$

where by consistent we mean that if $(b_i(s, \eta_i))_{i \in I}$ is the strategy profile in the Bayesian Nash equilibrium, then $\sigma_i(a_i, s) = \int_{\{\eta_i | b_i(s, \eta_i) = a_i\}} d\mu_i(\eta_i)$.⁸

⁸Given the absolute continuity of μ_i , player i has a unique best reply for almost all realizations of η_i .

This characterization of Bayesian Markov perfect equilibrium is similar in spirit to the characterization in Section 2.2. The main difference is that here we have a game of incomplete information and therefore the equilibrium concept in the reduced one-shot game is Bayesian Nash equilibrium. Escobar (2007) ensures the existence of a Bayesian Markov perfect equilibrium in the perturbed dynamic stochastic game.

6.2 Purification: Convergence and Approachability

We are now ready to explore how good of an approximation to the original (unperturbed) 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 to the equilibria of the original game.

Before answering this question, we provide a precise notion of convergence for a sequence of probability distributions.

Definition 13 *The sequence of probability distributions $(\mu_i^n)_{n \in \mathbb{N}}$ converges to a mass point at $0 \in \mathbb{R}^{|A|}$ as $n \rightarrow \infty$ if*

$$\lim_{n \rightarrow \infty} \frac{1}{\mu_i^n(R^n)} \int_{\eta_i \in R^n} \left\{ \max_{a \in A} |\eta_i(a)| \right\} d\mu_i^n(\eta_i) = 0$$

for any sequence of measurable sets $(R^n)_{n \in \mathbb{N}}$ such that $\mu_i^n(R^n) > 0$.

The perturbations considered by Harsanyi (1973a) satisfy Definition 13. While our approachability result in Theorem 15 remains valid under more general perturbations, we prefer to work with a single convergence notion because both convergence and approachability are desirable properties in applications. Note that Definition 13 is satisfied by, for example, any sequence of probability distribution $(\mu_i^n)_{n \in \mathbb{N}}$ such that the support of μ_i^n is contained in a ball of radius $r(n)$ centered at $0 \in \mathbb{R}^{|A|}$, where $r(n) \rightarrow 0$ as $n \rightarrow \infty$.

To facilitate 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 proposition shows that as the private shocks vanish, any converging sequence of equilibrium distributions for perturbed games converges to an equilibrium of the original game.

Proposition 14 (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 $n \rightarrow \infty$. Suppose further that $\sigma^n \in \mathbf{Equil}^n(U)$ converge to $\bar{\sigma}$ as $n \rightarrow \infty$. Then, $\bar{\sigma} \in \mathbf{Equil}(U)$.*

The proof of Proposition 14 is detailed in the Appendix.

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

Theorem 15 (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 $n \rightarrow \infty$. Let $\bar{\sigma}$ be a regular equilibrium of game U . Then, for all $\bar{\epsilon} > 0$ and all large enough n , there exists $\sigma^n \in \mathbf{Equil}^n(U)$ such that $\|\sigma^n - \bar{\sigma}\| < \bar{\epsilon}$.*

In conjunction with Theorem 5, Theorem 15 indicates that, for almost all games $U \in \mathbb{R}^{|I||A||S|}$, all equilibria are purifiable. Hence, one can interpret the original game as an idealization—a limit—of nearby games with a small amount of payoff uncertainty. Our main purification result also blunts a common criticism of the notion of a mixed-strategy equilibrium, namely that a player has no incentive to adhere to the prescribed randomization over his pure actions, since in the approximating equilibrium a player has a strictly optimal pure action for almost all realizations of his payoffs.

It is considerably more difficult to obtain lower hemi-continuity results such as Theorem 15 than closure results such as Proposition 14. The proof of Theorem 15 is detailed in Section 7.2 and relies on arguments previously presented by Govindan, Reny, and Robson (2003). That we are able to generalize their proof once again shows the power of our regularity notion. The two key properties satisfied by regular equilibria that we exploit are strong stability and quasi-strictness.

6.3 Application: Repeated Games

An prominent branch of research on repeated games focuses on how intertemporal incentives expand the range of equilibrium behaviors. In light of the well known folk theorem, the standard theory has little to say about what will happen: Any payoff vector dominating players' minimax payoffs can be supported as a subgame perfect equilibrium of the repeated game if players are patient enough. A natural question is thus whether introducing a small amount of incomplete information into the model is a sensible way to select some of these equilibria.

To answer this question, consider an arbitrary repeated game with a stage game that satisfies 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 Fudenberg and Maskin's (1986) folk theorem strategies constitute a subgame perfect equilibrium of the repeated game. Recall that these 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, followed by a carrot phase where the rivals are rewarded for carrying out the punishment phase. Moreover, after inspecting the proof of Theorem 2 in Fudenberg and Maskin (1986), it is easy to see that the incentives can be taken to be strict, both on- and off-path.

We proceed as in Example 4 to embed the repeated game and its folk theorem strategies into our dynamic stochastic game framework. We consider a dynamic stochastic game with a state space S having $1 + \sum_{i \in I} T_i + |I|$ elements. The first state amounts to the cooperation phase, the next $\sum_{i \in I} T_i$ states amount to the punishment phase, and the last $|I|$ states amount to the subsequent

carrot phase.⁹ The action sets and payoff functions of the dynamic stochastic game are the same as in the repeated game. The evolution of the play phase from one period to the next in the repeated game can be seen as a Markov process. We therefore define the transition function of the dynamic game to be such that, given the state s and the action profile a in the current period, the state s' in the subsequent period is the same as the play phase in the repeated game. The stationary Markov behavior strategies in the dynamic stochastic game mimic the folk theorem strategies in the obvious manner: In state s player i in the dynamic stochastic game plays as if he were playing in the repeated game in play phase s .

Since players in the repeated game have strict incentives to conform to the folk theorem strategies, both on- and off-path, the Markov perfect equilibrium in stationary Markov behavior strategies of the dynamic stochastic game is regular and thus purifiable. This implies that the subgame perfect equilibrium in folk theorem strategies is also purifiable.¹⁰ We conclude that introducing incomplete information into the repeated game cannot possibly yield a sensible selection procedure.

At the same time, our results allow us to conclude that another important class of subgame perfect equilibria in repeated games are knife-edge cases. Consider again the repeated prisoners' dilemma in Example 4 and introduce a payoff irrelevant state variable $s \in \{CC, CD, DC, DD\}$ with transition function

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

and initial distribution $\bar{q}(CC) = 1$. In words, the state in the subsequent period is the action profile played in the current period. Any strictly state dependent Markov perfect equilibrium of this dynamic stochastic game is a belief-free equilibrium of the repeated prisoners' dilemma. Introduced by Ely and Välimäki (2002), belief-free equilibria are a class of subgame perfect equilibria in stationary behavior strategies with one-period memory. Ely and Välimäki (2002) and Bhaskar, Mailath, and Morris (2007) show that there is a continuum of belief-free equilibria and that none of them is purifiable. These results are consistent with ours because, by restricting payoffs to be state independent, the repeated prisoners' dilemma constitutes a negligible subset of dynamic stochastic games. Corollary 7 implies that belief-free equilibria cannot be regular and thus cannot survive even a small amount of state dependence in payoffs.

⁹If some phase requires public randomization, then the state space can be expanded to represent the public randomization.

¹⁰We are implicitly assuming that, both on- and off-path, play is in pure strategies. This entails a loss of generality because it rules out mixed strategy minimax phases (see Theorem 5 in Fudenberg and Maskin 1986). Being nontrivially behavioral, these mixed strategies are not strict, and we have not been able to establish whether or not they are regular.

7 Proofs

In this section we detail the proofs of our main genericity and purification results in Theorems 5 and 15, respectively.

7.1 Proof of Theorem 5

7.1.1 Two Useful Lemmata

Below we present two lemmata. As a corollary to the second lemma we further obtain a characterization of the dimension of the equilibrium graph.

To facilitate the subsequent analysis we require some 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 with best replies B^* and carriers C^* . Formally,

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

where the sets $B_i(\sigma, s)$ and $C_i(\sigma, s)$ are as defined in Section 5.1. We also define $I(B^*, C^*)$ as the set of games having some irregular equilibrium with best replies B^* and carriers C^* . Clearly, $I(B^*, C^*) \subseteq G(B^*, C^*)$.

The first lemma shows that the set of games having some equilibrium that fails to be quasi-strict has measure zero. The proof proceeds as follows. We first derive a set of necessary conditions characterizing a game \bar{U} and an equilibrium $\bar{\sigma}$ that fails to be quasi-strict. These indifference conditions can be written as a system of equations, $M(\bar{\sigma}, \bar{U}) = 0$. Since these equations are linearly independent (as shown below in Claim 1), we can derive a locally defined function that maps strategies and *some* components of the payoff vector U to the *entire* vector U . We then show that the set of all games $G^{\bar{\sigma}, \bar{U}}(B^*, C^*)$ that are close to \bar{U} and have some equilibrium that is close to $\bar{\sigma}$ has a small dimension and is therefore negligible (Claim 2). The lemma finally follows by applying this logic to each possible pair $(\bar{\sigma}, \bar{U})$ in a properly chosen way.

Lemma 16 *If $B^* \neq C^*$, then $\lambda(G(B^*, C^*)) = 0$. It follows that $\lambda(I(B^*, C^*)) = 0$.*

Proof. Considering a game \bar{U} with an equilibrium $\bar{\sigma}$ such that $B_i(\bar{\sigma}, \cdot) = B_i^*$ and $C_i(\bar{\sigma}, \cdot) = C_i^*$ for all $i \in I$. Because $B^* \neq C^*$ by assumption, $\bar{\sigma}$ fails to be quasi-strict. Fix a collection of actions $a_i^s \in A_i$ such that $a_i^s \in C_i^*(s)$ for all $i \in I$ and $s \in S$.

By definition of $B_i^*(s)$, it must be that $\mathcal{U}_i(a_i, s, \bar{\sigma})$ equals $\mathcal{U}_i(a_i^s, s, \bar{\sigma})$ for all $a_i \in B_i^*(s)$. In words, player i 's payoff to all his best replies is the same. In matrix notation $\mathcal{U}_i(a_i, \bar{\sigma}, s)$ in equation (3.1) can be written as

$$\mathcal{P}_{a_i, \bar{\sigma}_{-i}}^s (\mathbb{I}_{|A||S|} - \delta_i Q \mathcal{P}_{\bar{\sigma}})^{-1} \bar{U}_i,$$

where $\mathcal{P}_{a_i, \bar{\sigma}_{-i}}^s$ is as defined in Section 2.3. Hence, for all $a_i \in B_i^*(s) \setminus \{a_i^s\}$,

$$(\mathcal{P}_{a_i, \bar{\sigma}_{-i}}^s - \mathcal{P}_{a_i^s, \bar{\sigma}_{-i}}^s)(\mathbb{I}_{|A||S|} - \delta_i Q \mathcal{P}_{\bar{\sigma}})^{-1} \bar{U}_i = 0.$$

For all $a_i \in B_i^*(s) \setminus \{a_i^s\}$ and $s \in S$, define the (a_i, s) row of $\bar{\mathcal{P}}_{i, \bar{\sigma}} \in \mathbb{R}^{\sum_{s \in S} (|B_i^*(s)| - 1) \times |A||S|}$ by $(\mathcal{P}_{a_i, \bar{\sigma}_{-i}}^s - \mathcal{P}_{a_i^s, \bar{\sigma}_{-i}}^s)$. Write the indifference conditions for player i as

$$\bar{\mathcal{P}}_{i, \bar{\sigma}}(\mathbb{I}_{|A||S|} - \delta_i Q \mathcal{P}_{\bar{\sigma}})^{-1} \bar{U}_i = 0.$$

Collect the indifference conditions for all players to obtain the system of equations

$$M(\bar{\sigma}, \bar{U}) = \begin{pmatrix} \bar{\mathcal{P}}_{1, \bar{\sigma}}(\mathbb{I}_{|A||S|} - \delta_1 Q \mathcal{P}_{\bar{\sigma}})^{-1} \bar{U}_1 \\ \vdots \\ \bar{\mathcal{P}}_{|I|, \bar{\sigma}}(\mathbb{I}_{|A||S|} - \delta_{|I|} Q \mathcal{P}_{\bar{\sigma}})^{-1} \bar{U}_{|I|} \end{pmatrix} = 0.$$

Claim 1 *The Jacobian $\frac{\partial M(\bar{\sigma}, \bar{U})}{\partial \bar{U}}$ has full rank $\sum_{i \in I} \sum_{s \in S} (|B_i^*(s)| - 1)$.*

The proof of Claim 1 is detailed in the Appendix. It follows from exploiting the diagonal structure of the Jacobian and by noting that $\bar{\mathcal{P}}_{i, \bar{\sigma}}$ has full rank and, as a consequence of Lemma 6, so does $\bar{\mathcal{P}}_{i, \bar{\sigma}}(\mathbb{I}_{|A||S|} - \delta_i Q \mathcal{P}_{\bar{\sigma}})^{-1}$.

As a consequence of the implicit function theorem, it is possible to obtain 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 $\mathcal{N} \subseteq \Sigma^\epsilon$, where $\bar{U} \in \mathcal{N}^1 \times \mathcal{N}^2$ (properly ordered) and $\bar{\sigma} \in \mathcal{N}$, and a function $\Phi: \mathcal{N}^1 \times \mathcal{N} \rightarrow \mathcal{N}^2$ such that for all $(\sigma, U^1) \in \mathcal{N} \times \mathcal{N}^1$, $\Phi(U^1, \mathcal{N}^1)$ is the unique solution $U^2 \in \mathcal{N}^2$ to $M(\sigma, (U^1, U^2)) = 0$. We define the function $H(\sigma, U^1) = (U^1, \Phi(\sigma, U^1))$. In order to highlight the dependence of these objects on $\bar{\sigma}$ and \bar{U} , we write $\mathcal{N}_{\bar{\sigma}, \bar{U}}^1$, $\mathcal{N}_{\bar{\sigma}, \bar{U}}^2$, $\mathcal{N}_{\bar{\sigma}, \bar{U}}$, and $H_{\bar{\sigma}, \bar{U}}$, respectively. We assume, without loss of generality, that $\mathcal{N}_{\bar{\sigma}, \bar{U}}^1$, $\mathcal{N}_{\bar{\sigma}, \bar{U}}^2$, and $\mathcal{N}_{\bar{\sigma}, \bar{U}}$ are balls with rational centers and radii.

Define $G^{\bar{\sigma}, \bar{U}}(B^*, C^*)$ as the set of all games that are close enough to \bar{U} and have some equilibrium that is close enough to $\bar{\sigma}$ with the same best replies and carriers as $\bar{\sigma}$. More formally,

$$G^{\bar{\sigma}, \bar{U}}(B^*, C^*) = \left\{ U \in \mathcal{N}_{\bar{\sigma}, \bar{U}}^1 \times \mathcal{N}_{\bar{\sigma}, \bar{U}}^2 \mid \text{there exists } \sigma \in \mathbf{Equil}(U) \cap \mathcal{N}_{\bar{\sigma}, \bar{U}} \right. \\ \left. \text{with } B_i(\sigma, \cdot) = B_i^* \text{ for all } i \in I \text{ and } \sigma \in A(C^*) \right\},$$

where $A(C^*) = \{\sigma \in \Sigma \mid C_i(\sigma, \cdot) = C_i^* \text{ for all } i \in I\}$. Further define the set

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

Clearly, $G^{\bar{\sigma}, \bar{U}}(B^*, C^*) \subseteq P^{\bar{\sigma}, \bar{U}}(B^*, C^*)$.

Claim 2 $\lambda(P^{\bar{\sigma}, \bar{U}}(B^*, C^*)) = 0$.

The proof of Claim 2 is as follows. Note that $\dim(\mathcal{N}_{\bar{\sigma}, \bar{U}}^1) = |I||A||S| - \sum_{i \in I} \sum_{s \in S} (|B_i^*(s)| - 1)$ and $\dim(A(C^*) \cap \mathcal{N}_{\bar{\sigma}, \bar{U}}) = \sum_{i \in I} \sum_{s \in S} (|C_i^*(s)| - 1)$. Therefore, $\dim(\mathcal{N}_{\bar{\sigma}, \bar{U}}^1 \times (A(C^*) \cap \mathcal{N}_{\bar{\sigma}, \bar{U}})) = |I||A||S| - \sum_{i \in I} \sum_{s \in S} |B_i^*(s)| + \sum_{i \in I} \sum_{s \in S} |C_i^*(s)| < |I||A||S|$ for $B_i^*(s) \neq C_i^*(s)$ for some $i \in I$ and $s \in S$. Since $P^{\bar{\sigma}, \bar{U}}(B^*, C^*) = H_{\bar{\sigma}, \bar{U}}((A(C^*) \cap \mathcal{N}_{\bar{\sigma}, \bar{U}}) \times \mathcal{N}_{\bar{\sigma}, \bar{U}}^1)$, the claim follows.

We are now ready to complete the proof of Lemma 16. For each 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 $i \in I$, we can construct the sets $G^{\bar{\sigma}, \bar{U}}(B^*, C^*)$ and $P^{\bar{\sigma}, \bar{U}}(B^*, C^*)$. Moreover, since the neighborhoods $\mathcal{N}_{\bar{\sigma}, \bar{U}}^1$, $\mathcal{N}_{\bar{\sigma}, \bar{U}}^2$, and $\mathcal{N}_{\bar{\sigma}, \bar{U}}$ are chosen from a countable set, it follows that

$$G(B^*, C^*) \subseteq \cup_{n \in \mathbb{N}} Q^n,$$

where $Q^n = P^{\bar{\sigma}^n, \bar{U}^n}(B^*, C^*)$ is constructed for each of the countable number of neighborhoods. Lemma 16 now follows from Claim 2 by noting that the union of measure zero sets has measure zero as well. ■

The proof of Lemma 16 resembles proofs given for normal form games by Harsanyi (1973a) and van Damme (1991). The main difference is that we do not define M globally. We indeed believe that this is impossible in our dynamic setting. We therefore analyze the system of equations $M(\bar{\sigma}, \bar{U}) = 0$ locally, and apply this construction to a countable set of games and equilibria.

Having disposed of all games having some equilibrium that fails to be quasi-strict, we turn to games having equilibria that all are quasi-strict. Within this class we restrict attention to completely mixed equilibria. The second lemma shows that for these equilibria the Jacobian of f with respect to the pair (σ, U) has full rank. Its proof is similar to that of Claim 1 and exploits Lemma 6 and the diagonal structure of the Jacobian.

To state the lemma, we define the set of completely mixed strategy profiles in Σ^ϵ as

$$\tilde{\Sigma} = \left\{ \sigma \in \Sigma^\epsilon \mid \sigma_i(a_i, s) > 0 \text{ for all } i \in I, a_i \in A_i, \text{ and } s \in S \right\}.$$

Lemma 17 *If $\sigma \in \tilde{\Sigma}$, then $\frac{\partial f(\sigma, U)}{\partial(\sigma, U)}$ has full rank $|S| \sum_{i \in I} |A_i|$.*

Proof. Considering a game U with a completely mixed equilibrium σ . Fix a collection of actions $a_i^s \in A_i$ for all $i \in I$ and $s \in S$.

In matrix notation $\sigma_i(a_i, s) \left(\mathcal{U}_i(a_i, s, \sigma) - \mathcal{U}_i(a_i^s, s, \sigma) \right)$ in equation (3.3) can be written as

$$\sigma_i(a_i, s) \left(\mathcal{P}_{a_i, \sigma_{-i}}^s - \mathcal{P}_{a_i^s, \sigma_{-i}}^s \right) \left(\mathbb{I}_{|A||S|} - \delta_i Q \mathcal{P}_\sigma \right)^{-1} U_i.$$

For all $a_i \in A_i \setminus \{a_i^s\}$ and $s \in S$, define the (a_i, s) row of $P_i^*(\sigma) \in \mathbb{R}^{(|A_i|-1)|S| \times |A||S|}$ by $\sigma_i(a_i, s) \left(\mathcal{P}_{a_i, \sigma_{-i}}^s - \right.$

$\mathcal{P}_{a_i^s, \sigma_{-i}}^s$). The components of f associated with player i can now be written as

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

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

$$\left(\begin{array}{ccc|ccc|c|ccc} \sigma_i(\cdot, s_1) & & & \sigma_i(\cdot, s_2) & & & \dots & & \sigma_i(\cdot, s_{|S|}) \\ 1 & \dots & 1 & 0 & \dots & 0 & \dots & 0 & \dots & 0 \\ 0 & \dots & 0 & 1 & \dots & 1 & \dots & 0 & \dots & 0 \\ \vdots & & \vdots & \vdots & & \vdots & & \vdots & & \vdots \\ 0 & \dots & 0 & 0 & \dots & 0 & \dots & 1 & \dots & 1 \end{array} \right).$$

This matrix, denoted by X_i , has full rank $|S|$. The Jacobian of the first $|S|$ components of f_i with respect to (σ_{-i}, U) is $0 \in \mathbb{R}^{|S| \times (\sum_{j \neq i} |A_j||S| + |I||A||S|)}$.

Next consider the components of f_i associated with $a_i \neq a_i^s$. The Jacobian of those components with respect to U takes the form

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

The matrix $P_i^*(\sigma)$ has full rank $|S|(|A_i| - 1)$. To see this, note that for all pairs (a_i, s) with $a_i \in A_i \setminus \{a_i^s\}$, the (a_i, s) row of $P_i^*(\sigma)$ 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 $(a_i, a_{-i}) = a'$). The (a_i, s) row also contains a nonzero term in some component (a', s) where a' contains a_i ; indeed, $\sum_{a' \text{ contains } a_i} \sigma_{-i}(a' \setminus a_i) = \sum_{a_{-i}} \sigma_{-i}(a_{-i}) = 1$, where given $a' = (a_i, a_{-i})$ we $a' \setminus a_i = a_{-i}$. Since $P_i^*(\sigma)$ has full rank, so does the matrix $Z_i = P_i^*(\sigma)(\mathbb{I}_{|A||S|} - \delta_i Q \mathcal{P}_\sigma)^{-1}$ as a consequence of Lemma 6.

We now see that, up to permutations of rows, the Jacobian of f with respect to the pair (σ, U)

takes the form

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

where $Y_i \in \mathbb{R}^{|\mathcal{S}| \sum_{i \in I} (|A_i| - 1) \times |A_i| |\mathcal{S}|}$. This permits us to deduce that $\frac{\partial f(\sigma, U)}{\partial(\sigma, U)}$ has full rank $|\mathcal{S}| \sum_{i \in I} |A_i|$.

■

Lemma 17 implies a version of the structure theorem.

Corollary 18 *The equilibrium graph has the same dimension as the space of games.*

Corollary 18 generalizes an observation made by Govindan and Wilson (2001) for normal form games to dynamic stochastic games. Had we only been interested in obtaining a finiteness result, this corollary and the transversality theorem yield it almost immediately. (To obtain a finiteness result, it is enough to restrict attention to completely mixed equilibria; therefore, there is no need for Lemma 16.) Our notion of regularity therefore leads to a much simpler finiteness proof than that of Haller and Lagunoff (2000). Our analysis also provides a finiteness proof for normal form games that cannot be deemed as more complicated than the proof based on the theory of semi-algebraic sets in Govindan and Wilson (2001).¹¹

7.1.2 Proof of Theorem 5

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

Theorem 19 (Transversality Theorem) *Let $\mathcal{O} \subseteq \mathbb{R}^n$ be open and $L: \mathcal{O} \times \mathbb{R}^s \rightarrow \mathbb{R}^n$ be continuously differentiable. Assume that the Jacobian $\frac{\partial L(x, y)}{\partial(x, y)}$ 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, for almost all $\bar{y} \in \mathbb{R}^s$, 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 Mas-Colell, Whinston, and Green (1995) for an intuitive discussion and applications in economics and Abraham and Robbin (1967) and Guillemin and Pollack (1974) for further results and technical details.

¹¹Note, however, that extending Govindan and Wilson's (2001) tools to dynamic stochastic games requires establishing the semi-algebraicity of f .

We are now ready to prove Theorem 5. 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 only a finite number of sets B^* and C^* such that $C^* \subseteq B^*$, it is enough to show that $\lambda(I(B^*, C^*)) = 0$ for all such sets. If $B^* \neq C^*$, this follows from Lemma 16.

If $B^* = C^*$, then equation (5.1) in the proof of Proposition 8 implies that $J(\sigma) = \frac{\partial f(\sigma, U)}{\partial \sigma}$ has full rank if and only if so does the submatrix $\bar{J}(\sigma)$ obtained from $J(\sigma)$ by crossing out all rows and columns corresponding to components (a_i, s) with $a_i \notin B_i^*(s)$. This submatrix is itself the Jacobian of a completely mixed equilibrium. Without loss of generality we can therefore assume that $B^*(s) = C^*(s) = A$ for all $s \in S$. Then,

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

From Lemma 17, $\frac{\partial f(\sigma, U)}{\partial(\sigma, U)}$ has full rank for all pairs $(\sigma, U) \in \tilde{\Sigma} \times \mathbb{R}^{|I||A||S|}$. The transversality theorem therefore implies that $\frac{\partial f(\sigma, U)}{\partial \sigma}$ has full rank for almost all games $U \in \mathbb{R}^{|I||A||S|}$.

7.2 Proof of Theorem 15

To prove Theorem 15 we proceed as follows. In Section 7.2.1 we first derive a system of nonlinear equations that characterizes the equilibrium distributions of a perturbed dynamic stochastic game. In Section 7.2.2 we then exploit a result from algebraic topology to ensure that there exists a solution to this system and, moreover, that this solution is close enough to the regular equilibrium $\bar{\sigma}$ of the original (unperturbed) game.

7.2.1 Alternative Characterization

Below we derive a system of nonlinear equations that characterizes the equilibrium distributions of a perturbed dynamic stochastic game. This, in effect, amounts to providing an alternative characterization of a Bayesian Markov perfect equilibrium. See Aguirregabiria and Mira (2007) for a similar derivation.

Continuation Values. Consider a dynamic stochastic game with perturbations $(\mu_i)_{i \in I}$ and equilibrium strategy profile b . Let σ^b be the corresponding consistent distribution profile. Then $\bar{V}_i: S \rightarrow \mathbb{R}$, the equilibrium value function for player i , is the solution the Bellman equation

$$\bar{V}_i(s) = u_i(\sigma^b(\cdot, s), s) + \sum_{a_i \in A_i} \int_{\{\eta_i \mid b_i(s, \eta_i) = a_i\}} \eta_i(a_i, \sigma_{-i}^b(\cdot, s)) d\mu_i(\eta_i) + \delta_i \sum_{s' \in S} \bar{V}_i(s') q(s'; \sigma^b(\cdot, s), s). \quad (7.1)$$

The first and the third term on the right hand side of equation (7.1) depend on b only indirectly through σ^b . Proposition 1 in Hotz and Miller (1993) ensures that

$$\sum_{a_i \in A_i} \int_{\{\eta_i | b_i(s, \eta_i) = a_i\}} \eta_i(a_i, \sigma_{-i}^b(\cdot, s)) d\mu_i(\eta_i) = e_i(\sigma^b, s),$$

where $e_i(\sigma^b, s)$ is the expected value of the private shock given optimizing behavior. Hence, the second term in equation (7.1) is seen to also depend b only indirectly through σ^b . See Aguirregabiria and Mira (2002) for further discussion.

Importantly, $e_i(\sigma, s)$ is well defined even if $\sigma \in \Sigma$ is not an equilibrium distribution. Moreover, $e_i(\sigma, s)$ is a continuous function of $\sigma \in \Sigma$. We note that for all $s \in S$ the range of $e_i(\cdot, s): \Sigma \rightarrow \mathbb{R}$ is contained in the interval $[-\gamma_i, \gamma_i]$, where $\gamma_i = |A_i| \int \max_{a \in A} \{|\eta_i(a)|\} d\mu_i(\eta_i)$. According to Tietze's extension theorem (Royden 1968), it is therefore possible to extend $e_i(\cdot, s)$ to $\bar{\Sigma}^\epsilon$, the closure of Σ^ϵ , in a continuous manner such that its range is contained in the interval $[-\gamma_i, \gamma_i]$. Slightly abusing notation, we denote the extended function by $e_i(\cdot, s): \bar{\Sigma}^\epsilon \rightarrow \mathbb{R}$.

Using this construction, the value of continued play given an arbitrary profile $\sigma \in \bar{\Sigma}^\epsilon$ is

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

where the s component of $e_i(\sigma) \in \mathbb{R}^{|S|}$ is given by $e_i(\sigma, s)$. We interpret $\bar{V}_i(s, \sigma)$ as the expected net present value of the stream of payoffs to player i if the dynamic system is currently in state s and play is according to σ . Note that the above formula above reduces to equation (2.3) if $\mu_i(\{0\}) = 1$.

Equilibrium Distributions. Fix $\sigma \in \bar{\Sigma}^\epsilon$ and let $\bar{V}_i(\cdot, \sigma)$ be the corresponding value of continuation play. Define the best reply of player i in state s as

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

$b_i^\sigma(s, \eta^i)$ is the best reply of player i if the current state is s , his private shock is η_i , his rivals play according σ_{-i} , and player i plays according to σ_i from the subsequent period on.

Define the (i, a_i, s) component of the function $g: \bar{\Sigma}^\epsilon \rightarrow \Sigma$ by

$$g_{i, a_i, s}(\sigma) = \int_{\{\eta_i | b_i^\sigma(s, \eta_i) = a_i\}} d\mu_i(\eta_i). \quad (7.2)$$

$g_{i, a_i, s}(\sigma)$ is the probability that the best reply of player i in state s is a_i . The following lemma characterizes the equilibrium distributions of the dynamic stochastic game with perturbations $(\mu_i)_{i \in I}$ as fixed points of g .

Lemma 20 *A profile $\sigma \in \bar{\Sigma}^\epsilon$ is an equilibrium distribution if and only if $g(\sigma) = \sigma$.*

This lemma is standard up to the fact that the domain of g is not Σ but $\bar{\Sigma}^\epsilon$. It follows because the range of g is contained in Σ so that a fixed point of g must belong to Σ .

Finally, define the (i, a_i, s) component of the function $h: \bar{\Sigma}^\epsilon \rightarrow \mathbb{R}^{|\mathcal{S}| \sum_{i \in I} |A_i|}$ by

$$h_{i,a_i,s}(\sigma) = \begin{cases} \sum_{a_i \in A_i} \sigma_i(a_i, s) - 1 & \text{if } a_i = a_i^s, \\ g_{i,a_i,s}(\sigma) - \sigma_i(a_i, s) & \text{if } a_i \neq a_i^s, \end{cases}$$

where a_i^s is the reference action for player i in state s as used in the construction of the function of f for the equilibrium $\bar{\sigma}$ of the unperturbed game U . Since g is continuous in σ , so is h . It is not hard to see that $h(\sigma) = 0$ if and only if $g(\sigma) = \sigma$ so the problem of finding an equilibrium distribution reduces to finding a zero of h .

7.2.2 Proof of Theorem 15

We employ the following result from algebraic topology.

Proposition 21 (Govindan, Reny, and Robson (2003)) *Suppose that O is a bounded, open set in \mathbb{R}^m and $h, f: \bar{O} \rightarrow \mathbb{R}^m$ 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 all $t \in [0, 1]$, the function $th + (1-t)f$ has no zero on the boundary of O , then h has a zero in O .*

The equilibrium $\bar{\sigma}$ of the unperturbed game U is a zero of f . Also we have constructed h so that a zero of h is an equilibrium distribution of the perturbed game. In what follows we use Proposition 21 to establish the existence of a Bayesian Markov perfect equilibrium of the perturbed game.

Since the equilibrium $\bar{\sigma}$ of the unperturbed game U is regular, the argument developed in the proof of Proposition 9 shows that there exists an open set $O \subseteq \bar{\Sigma}^\epsilon$ that satisfies the following conditions (referred to hereafter as C1-C5):

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

C2 $\|\bar{\sigma} - \sigma\| < \bar{\epsilon}$ for all $\sigma \in O$;

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

C4 for all $i \in I$, $a_i \in A_i$, and $s \in S$, if $\bar{\sigma}_i(a_i, s) > 0$, then $\sigma_i(a_i, s) > 0$ for all $\sigma \in O$;

C5 for all $i \in I$, $a_i \in A_i$, and $s \in S$, if $\mathcal{U}_i(a_i, s, \bar{\sigma}) - \mathcal{U}_i(a_i^s, s, \bar{\sigma}) < 0$, then $\mathcal{U}_i(a_i, s, \sigma) - \mathcal{U}_i(a_i^s, s, \sigma) < 0$ for all $\sigma \in O$.

Consider the sequence of probability distributions of private shocks $(\mu_i^n)_{n \in \mathbb{N}}$. For all $i \in I$, use μ_i^n to construct e_i^n , \bar{V}_i^n , g^n , and h^n as detailed in Section 7.2.1. To prove Theorem 15 it suffices

to find a zero of h^n in O for all large enough n . Such a zero is an equilibrium distribution of the dynamic stochastic game with perturbations $(\mu_i^n)_{i \in I}$ and it is within a distance at most $\bar{\epsilon}$ of $\bar{\sigma}$ due to C2. As a consequence of Proposition 21, the following lemma yields the desired result.

Lemma 22 *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. 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 $t^n h^n + (1 - t^n)f(\cdot, U)$. We state and prove three preliminary claims.

Claim 3 *If $\mathcal{U}_i(a_i, s, \hat{\sigma}) - \mathcal{U}_i(a_i^s, s, \hat{\sigma}) < 0$, then $g_{i, a_i, s}^n(\sigma^n) \rightarrow 0$.*

The proof of this claim is as follows. For $a'_i \in A_i$, $s \in S$, and $\sigma \in \Sigma^\epsilon$, define

$$\mathcal{U}_i^n(a'_i, s, \sigma) = u_i(a'_i, \sigma_{-i}(\cdot, s), s) + \delta_i \sum_{s' \in S} \bar{V}_i^n(s', \sigma) q(s'; a'_i, \sigma_{-i}(\cdot, s), s). \quad (7.3)$$

Note that

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

because the range of $e_i^n(\cdot, s)$ is contained in $[-\gamma_i^n, \gamma_i^n]$, where

$$\gamma^n = |A_i| \int \left\{ \max_{a \in A} |\eta_i(a)| \right\} d\mu_i^n(\eta_i) \rightarrow 0.$$

It thus follows that $\mathcal{U}_i^n(a'_i, s, \sigma^n) \rightarrow \mathcal{U}_i(a'_i, s, \hat{\sigma})$ for all $a'_i \in A_i$ and $s \in S$. Consequently, there exists $\psi > 0$ such that for all large enough n

$$\mathcal{U}_i^n(a_i, s, \sigma^n) - \mathcal{U}_i^n(a_i^s, s, \sigma^n) < -\psi.$$

By the definition of g in equation (7.2),

$$\begin{aligned} g_{i, a_i, s}^n(\sigma^n) &\leq \int_{\{\eta_i \in \mathbb{R}^{|A|} \mid |\eta_i(a_i, \sigma_{-i}^n(\cdot, s)) - \eta_i(a_i^s, \sigma_{-i}^n(\cdot, s))| \geq \mathcal{U}_i^n(a_i^s, s, \sigma^n) - \mathcal{U}_i^n(a_i, s, \sigma^n)\}} d\mu_i^n(\eta_i) \\ &\leq \int_{\{\eta_i \in \mathbb{R}^{|A|} \mid |\eta_i(a_i, \sigma_{-i}^n(\cdot, s)) - \eta_i(a_i^s, \sigma_{-i}^n(\cdot, s))| \geq \psi\}} d\mu_i^n(\eta_i) \\ &\leq \int_{\{\eta_i \in \mathbb{R}^{|A|} \mid |\eta_i(a_i, \sigma_{-i}^n(\cdot, s)) - \eta_i(a_i^s, \sigma_{-i}^n(\cdot, s))| \geq \psi\}} \frac{|\eta_i(a_i, \sigma_{-i}^n(\cdot, s)) - \eta_i(a_i^s, \sigma_{-i}^n(\cdot, s))|}{\psi} d\mu_i^n(\eta_i) \\ &\leq \frac{2}{\psi} \int \left\{ \max_{a \in A} |\eta_i(a)| \right\} d\mu_i^n(\eta_i) \\ &\rightarrow 0. \end{aligned}$$

This completes the proof of Claim 3.

Claim 4 *If $\mathcal{U}_i(a_i, s, \hat{\sigma}) - \mathcal{U}_i(a_i^s, s, \hat{\sigma}) < 0$, then $\hat{\sigma}_i(a_i, s) = 0$.*

The proof of this claim is as follows. From Claim 3, $g_{i,a_i,s}^n(\sigma^n) \rightarrow 0$. Therefore,

$$\begin{aligned} 0 &= \lim_{n \rightarrow \infty} t^n h_{i,a_i,s}^n(\sigma^n) + (1 - t^n) f_{i,a_i,s}(\sigma^n, U) \\ &= \hat{t}(-\hat{\sigma}_i(a_i, s)) + (1 - \hat{t}) f_{i,a_i,s}(\hat{\sigma}, U) \\ &= -\hat{\sigma}_i(a_i, s) \left(\hat{t} + (1 - \hat{t})(\mathcal{U}_i(a_i^s, s, \hat{\sigma}) - \mathcal{U}_i(a_i, s, \hat{\sigma})) \right). \end{aligned}$$

It follows that $\hat{\sigma}_i(a_i, s) = 0$. This completes the proof of Claim 4.

Claim 5 *For all $i \in I$ and $s \in S$, $\hat{\sigma}_i(\cdot, s)$ is a probability distribution. In addition, there exist $i \in I$, $a_i \in A_i \setminus \{a_i^s\}$, and $s \in S$ such that $f_{i,a_i,s}(\hat{\sigma}, U) \neq 0$.*

The proof of this claim is as follows. If $\hat{\sigma}_i(a_i, s) < 0$, then C4 implies that $\bar{\sigma}_i(a_i, s) = 0$. Since $\bar{\sigma}$ is quasi-strict, $\mathcal{U}_i(a_i, s, \bar{\sigma}) - \mathcal{U}_i(a_i^s, s, \bar{\sigma}) < 0$ and, from C5, $\mathcal{U}_i(a_i, s, \hat{\sigma}) - \mathcal{U}_i(a_i^s, s, \hat{\sigma}) < 0$. Claim 4 shows that $\hat{\sigma}_i(a_i, s) = 0$, a contradiction. Therefore, $\hat{\sigma}_i(a_i, s) \geq 0$ for all $i \in I$, $a_i \in A_i$, and $s \in S$. Further, because $h_{i,a_i^s,s}^n \equiv f_{i,a_i^s,s}$, 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 over actions. From C3, $\hat{\sigma}$ cannot be a zero of $f(\cdot, U)$. So there must exist $i \in I$, $a_i \in A_i \setminus \{a_i^s\}$, and $s \in S$ such that $f_{i,a_i,s}(\hat{\sigma}, U) \neq 0$. This completes the proof of Claim 5.

With these claims in hand, we are ready to complete the proof of Lemma 22. Fix $i \in I$, $a_i \in A_i \setminus \{a_i^s\}$, and $s \in S$ as in Claim 5 for the rest of the proof. Note that a_i^s cannot belong to $B_i(\hat{\sigma}, s)$. Indeed, if it did, then $\mathcal{U}_i(a_i, s, \hat{\sigma}) - \mathcal{U}_i(a_i^s, s, \hat{\sigma}) \leq 0$ and, from Claim 4, $f_{i,a_i,s}(\hat{\sigma}, U) = 0$, contradicting the definition of a_i (Claim 5). Since $\bar{\sigma}$ is quasi-strict, $\bar{\sigma}_i(a_i', s) > 0$ for all $a_i' \in B_i(\bar{\sigma}, s)$; this together with C4 implies that $\hat{\sigma}_i(a_i', s) > 0$ 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,

$$\sum_{a_i' \in B_i(\hat{\sigma}, s)} f_{i,a_i',s}(\sigma^n, U) > 0. \quad (7.4)$$

For $a_i' \notin B_i(\hat{\sigma}, s)$, $g_{i,a_i',s}^n(\sigma^n) \rightarrow 0$, so that

$$\sum_{a_i' \in B_i(\hat{\sigma}, s)} g_{i,a_i',s}^n(\sigma^n) \rightarrow 1.$$

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,

$$\sum_{a_i' \in B_i(\hat{\sigma}, s)} h_{i,a_i',s}^n(\sigma^n) > 0 \quad (7.5)$$

for large enough n . But equations (7.4) and (7.5) imply that $t^n h^n + (1 - t^n)f(\cdot, U)$ is not zero at σ^n , a contradiction. This completes the proof of Lemma 22. ■

A Appendix

In this Appendix we detail the construction of Σ^ϵ and the proofs of some supporting results.

A.1 Construction of Σ^ϵ

We construct Σ^ϵ is as follows. Note that, for all $\sigma \in \Sigma$, $\mathbb{I}_{|S|} - \delta_i \mathcal{P}_\sigma Q$ and $\mathbb{I}_{|A||S|} - \delta_i Q \mathcal{P}_\sigma$ are invertible. Indeed, $\mathcal{P}_\sigma Q$ and $Q \mathcal{P}_\sigma$ are stochastic matrices so that $\mathbb{I}_{|S|} - \delta_i \mathcal{P}_\sigma Q$ and $\mathbb{I}_{|A||S|} - \delta_i Q \mathcal{P}_\sigma$ have strictly dominant diagonals. Therefore, for all $\bar{\sigma} \in \Sigma$, we can find $\epsilon_{\bar{\sigma}} > 0$ such that $\mathbb{I}_{|S|} - \delta_i \mathcal{P}_\sigma Q$ and $\mathbb{I}_{|A||S|} - \delta_i Q \mathcal{P}_\sigma$ are invertible for all $\sigma \in \mathbb{R}^{\sum_{i \in I} |A_i||S|}$ satisfying $\|\bar{\sigma} - \sigma\| < \epsilon_{\bar{\sigma}}$. Since Σ is compact, we can take a finite covering $(B(\bar{\sigma}^j, \epsilon_{\bar{\sigma}^j}))_{j \in J}$ of Σ . Define Σ^ϵ to be open such that its closure, denoted $\bar{\Sigma}^\epsilon$, is contained in the open set $\cup_{j \in J} B(\bar{\sigma}^j, \epsilon_{\bar{\sigma}^j})$.

A.2 Omitted Proofs

Proof of Claim 1. The Jacobian of M with respect to U takes the form

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

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 all pairs (a_i, s) with $a_i \in B_i^*(s) \setminus \{a_i^s\}$, the (a_i, s) row of $\bar{\mathcal{P}}_{i, \bar{\sigma}}$ 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 $(a_i, a_{-i}) = a'$). The (a_i, s) row also contains a nonzero term in some component (a', s) where a' contains a_i ; indeed, $\sum_{a' \text{ contains } a_i} \bar{\sigma}_{-i}(a' \setminus a_i, s) = \sum_{a_{-i}} \bar{\sigma}_{-i}(a_{-i}, s) = 1$, where given $a' = (a_i, a_{-i})$ we take $a' \setminus a_i = a_{-i}$. This shows that $\bar{\mathcal{P}}_{i, \bar{\sigma}}$ has full rank.

The matrix $(\mathbb{I}_{|A||S|} - \delta_i Q \mathcal{P}_{\bar{\sigma}})^{-1}$ has full rank as a consequence of Lemma 6. Taken together, this and the above observation imply that the Jacobian of M with respect to U has full rank $\sum_{i \in I} \sum_{s \in S} (|B_i^*(s)| - 1)$. ■

Proof of Proposition 14. For large enough n , $C_i(\bar{\sigma}, s) \subseteq C_i(\sigma^n, s)$ for all $i \in I$ and $s \in S$. By definition of g in equation (7.2), for any $a_i \in C_i(\sigma^n, s)$, $g_{i, a_i, s}^n(\sigma^n) = \sigma_i^n(a_i, s) > 0$. Therefore, there exists a set $R_{i, s}^n \subseteq \mathbb{R}^{|A|}$ such that $\mu_i^n(R_{i, s}^n) > 0$ and, for all $\eta_i \in R_{i, s}^n$ and $a'_i \in A_i$,

$$\mathcal{U}_i^n(a_i, s, \sigma^n) - \mathcal{U}_i^n(a'_i, s, \sigma^n) > \eta_i(a'_i, \sigma_{-i}^n(\cdot, s)) - \eta_i(a_i, \sigma_{-i}^n(\cdot, s)),$$

where \mathcal{U}_i^n is defined in equation (7.3). We can integrate out this inequality to deduce that, for all $a'_i \in A_i$,

$$\mathcal{U}_i^n(a_i, s, \sigma^n) - \mathcal{U}_i^n(a'_i, s, \sigma^n) > \frac{1}{\mu_i^n(R_{i,s}^n)} \int_{\eta_i \in R_{i,s}^n} \eta_i(a'_i, \sigma_{-i}^n(\cdot, s)) - \eta_i(a_i, \sigma_{-i}^n(\cdot, s)) d\mu_i^n(\eta_i).$$

Letting $n \rightarrow \infty$ it follows that, for all $a'_i \in A_i$,

$$\mathcal{U}_i(a_i, s, \sigma) - \mathcal{U}_i(a'_i, s, \sigma) \geq 0.$$

We have therefore shown that for any $a_i \in C_i(\bar{\sigma}, s)$, $a_i \in B_i(\bar{\sigma}, s)$. This proves the result. ■

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 C. HO (2007): “A dynamic oligopoly game of the US airline Industry: Estimation and policy experiments,” Working paper, University of Toronto, Toronto.
- AGUIRREGABIRIA, V., AND P. MIRA (2002): “Swapping the nested fixed point algorithm: A class of estimators for discrete Markov decision models,” *Econometrica*, 70(4), 1519–1543.
- (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.
- BENKARD, L. (2004): “A dynamic analysis of the market for wide-bodied commercial aircraft,” *Review of Economic Studies*, 71(3), 581–611.
- 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.
- BESANKO, D., U. DORASZELSKI, Y. KRYUKOV, AND M. SATTERTHWAITTE (2007): “Learning-by-doing, organizational forgetting, and industry dynamics,” Working paper, Harvard University, Cambridge.
- BHASKAR, V. (1998): “Informational constraints and the overlapping generations model: Folk and anti-folk theorems,” *Review of Economic Studies*, 65(1), 135–149.
- (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,” Working paper no. 07-024, Penn Institute for Economic Research, Philadelphia.

- BORKOVSKY, R., U. DORASZELSKI, AND S. KRYUKOV (2008): “A user’s guide to solving dynamic stochastic games using the homotopy method,” Working paper, Northwestern University, Evanston.
- COLLARD-WEXLER, A. (2005): “Plant turnover and demand fluctuations in the ready-mix concrete industry,” Working paper, New York University, New York.
- 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. SATTERTHWAITTE (2007): “Computable Markov-perfect industry dynamics: Existence, purification, and multiplicity,” Working paper, Harvard University, Cambridge.
- DUBE, J., G. HITSCH, AND P. MANCHANDA (2005): “An empirical model of advertising dynamics,” *Quantitative Marketing and Economics*, 3, 107–144.
- ELY, J., AND J. VÄLIMÄKI (2002): “A robust folk theorem for the prisoners dilemma,” *Journal of Economic Theory*, 102(1), 84–105.
- ERDEM, E., AND J. TYBOUT (2003): “Trade policy and industrial sector responses: 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.
- 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.
- GOWRISANKARAN, G., AND R. TOWN (1997): “Dynamic equilibrium in the hospital industry,” *Journal of Economics and Management Strategy*, 6(1), 45–74.
- 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.
- HOTZ, J., AND R. MILLER (1993): “Conditional choice probabilities and the estimation of dynamic models,” *Review of Economic Studies*, 60(3), 497–529.
- JUDD, K., AND C. SU (2008): “Constrained optimization approaches to estimation of structural models,” Working paper, Hoover Institution, Stanford.
- 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 Research*, 10, 650–663.
- MAS-COLELL, A., M. WHINSTON, AND J. GREEN (1995): *Microeconomic theory*. Oxford University Press, Oxford.
- MASKIN, E., AND J. TIROLE (2001): “Markov perfect equilibrium, I: Observable actions,” *Journal of Economic Theory*, 100, 191–219.
- PAKES, A. (2008): “Theory and empirical work on imperfectly competitive markets,” Fisher-Schultz lecture, Harvard University, Cambridge.
- 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 (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 (2008): “Asymptotic least squares estimators for dynamic games,” *Review of Economic Studies*, forthcoming.
- ROYDEN, H. (1968): *Real analysis*. Macmillan, New York.
- RYAN, S. (2005): “The costs of environmental regulation in a concentrated industry,” Working paper, Duke University, Durham.
- 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.
- 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.
- ZANGWILL, W., AND C. GARCIA (1981): *Pathways to solutions, fixed points, and equilibria*. Prentice Hall, Englewood Cliffs.