

Genericity and Markovian Behavior in Stochastic Games

Hans Haller* and Roger Lagunoff†

Revised: May 29, 1999‡

Abstract

This paper examines *Markov Perfect equilibria* of general, finite state stochastic games. Our main result is that the number of such equilibria is finite for a set of stochastic game payoffs with full Lebesgue measure. We further discuss extensions to lower dimensional stochastic games like the alternating move game.

Journal of Economic Literature Classification Numbers: C72, C73.

Keywords: stochastic games, Markov Perfect equilibria, genericity.

*Department of Economics, Virginia Polytechnic Institute and State University, Blacksburg, VA 24061 USA. Phone: (540) 231-7591. e-mail: *haller@vt.edu*

†Department of Economics, Georgetown University, Washington DC 20057 USA. Phone: (202) 687-1510. e-mail: *lagunofr@gunet.georgetown.edu*.

‡We would like to thank, without implicating, Prajit Dutta, Daniel Kahn, Andrew McLennan, Lones Smith, and especially Akihiko Matsui for helpful conversations. We are also grateful to Drew Fudenberg and three anonymous referees for very helpful and detailed comments and suggestions.

1 Introduction

If a finite action stage game is infinitely repeated, then the standard theory of repeated games has little to say about what will happen. It is well known from the Folk Theorem that payoffs which dominate the minmax payoff can be supported as Perfect equilibria if players are patient enough.¹ Moreover, the indeterminacy problem extends to more general models of intertemporal play. Dutta (1995) proves a Folk Theorem for *stochastic games* satisfying a full dimensionality condition. Stochastic games, introduced by Shapley (1953), are general models of intertemporal play in which the stage game evolves according to a transition function which maps (probabilistically) the current state and action profile into next period's state.

One approach taken to address the indeterminacy problem considers strategic behavior of finite complexity. Various forms of finite automata, for example, are shown to reduce multiplicity in some repeated games.² One natural class of simple automata receiving considerable attention is the class of Markovian strategies. Markovian strategies rule out the possibility to extend memory, by using states of the automaton, beyond what is encoded in the states of the game. Hence, players condition only on payoff-relevant information.³ The present paper establishes a genericity result with respect to Markov Perfect equilibria (MPE) in stochastic games. Specifically, we show that the set of Markov Perfect equilibria is finite for a set of stochastic game payoffs with full Lebesgue measure.

In the special case where the stochastic game is a standard repeated game, this result is not surprising. Since players' Markov strategies do not condition on the past, the Markov Perfect equilibria of repeated games consist of infinite repetitions of Nash equilibria of the stage game. In that case, one can apply the well known result of Wilson (1971) and Harsanyi (1973) who show that the number of Nash equilibria in any normal form game is generically finite with respect to payoffs.⁴ Hence, in a trivial sense, Markov Perfect equilibria break the extreme multiplicity in repeated games by preventing all intertemporal reactions.

¹See Aumann (1981) and Fudenberg and Maskin (1986).

²See, for example, Abreu and Rubinstein (1988), Kalai, Samet and Stanford (1988), and Cho and Li (1999), among others. The aforementioned work first fixes bounds on states, then examines outcomes arising from them. Without fixing the bounds *a priori*, Kalai and Stanford (1988) show that in the standard repeated game any perfect equilibrium can be approximated by strategies which are finitely complex in the sense that they are implementable by finite-state automata.

³Recent work focussing on Markovian behavior includes Dutta (1996), Maskin and Tirole (1997), and Fudenberg and Maskin (1995).

⁴Other analyses of genericity in games include Rosenmüller (1971), Kreps and Wilson (1982), Blume and Zame (1994), Park (1997), Govidan and McLennan (1998a), Keiding (1997), McKelvey and McLennan (1997), McLennan (1997), Quint and Shubik (1997), McLennan and Park (1999). The last five derive generic upper bounds for the number of equilibria.

However, the Harsanyi-Rosenmüller-Wilson result cannot be directly applied to general stochastic games. Unlike in normal form games, players' payoffs span an infinite horizon and are not linear in strategies. Nevertheless, with some effort, the recursive structure of Markovian behavior admits a proof of generic finiteness similar in spirit to techniques used by Debreu (1970).

Section 2 gives a brief account of various notions of genericity. While “generic finiteness” does not eliminate multiplicity, it demonstrates that MPE are locally isolated, and so comparative statics exercises for small changes are well defined. Section 3 introduces definitions and notation. Section 4 describes the main result. There, we examine the logic of the result, and give a formal proof. Besides the unsurprising case of repeated games, our result applies to many nontrivial stochastic games of interest. Examples include resource extraction games (e.g., Levhari and Mirman (1980)), capital accumulation games (e.g., Dutta and Sundaram (1992)), and strategic learning models (e.g., Bergemann and Valimaki (1996)).

One particular stochastic game to which our result does *not* apply is the alternating move game. In these games, two players move in sequence; each player can only revise his decision in alternate periods. Examples are the duopoly price and Cournot quantity games studied in Benabou (1989) and Maskin and Tirole (1987, 1988a,b). As it turns out, these alternating move games are nongeneric in the fully dimensional state-dependent payoff space. Section 5 discusses reasons for this and examines potential extensions to lower dimensional games. A companion paper, Haller and Lagunoff (1997), in fact, extends the genericity result to the alternating move game. Section 6 contains an Appendix with details of the proof.

Note that unlike in repeated games, the Markov restriction in general stochastic games may encode enough information to construct useful punishments and support cooperation. For instance, MPE may be shown to support cooperation in certain state-dependent versions of repeated Prisoner's Dilemma. One example is the alternating move Prisoner's Dilemma game. Each player can alter his behavior only in every other period. It is not difficult to show that mixed strategy punishments may be used to enforce cooperation when players are patient enough. Alternatively, a slight modification of the standard repeated Prisoner's Dilemma game adds a state variable which can support cooperation. In this modification, a nontrivial state variable varies depending on whether or not a defection has been observed. The stage payoffs can then be made to vary across states by adding a very small number to each of the stage game payoffs in state D (the state in which someone defected). This perturbed game now has a payoff relevant state variable which can clearly be used to trigger an infinite “Nash-threats” punishment should a defection ever occur.

In other types of stochastic games Dutta (1996) establishes existence of efficient MPE. His result utilizes a sufficient condition called “state symmetry” which requires that all in-

dividuals essentially rank the states in the same way (in their preferences) in the symmetric optimum. For example, common pool resource extraction problems satisfy state symmetry.⁵ The intuition is that efficiency is attainable when players are patient since individuals' ranking over continuation payoffs coincide.

2 Genericity Concepts

We have argued that choosing the appropriate parameter space is crucial. For this reason the choice of genericity concept also warrants attention. There are two notions of genericity: one is measure-theoretic, and the other is topological. Our concept of genericity is full Lebesgue measure, the standard measure-theoretical notion of genericity in a Euclidean space or a suitable subset thereof. The term “generic” is meant to capture the fact that a property holds “for almost all parameters.” More formally, a property is called generic for a family of models indexed by parameters $b \in B$ if the property holds for a generic subset A of B . The standard notion of genericity in the measure-theoretical sense assumes that B is a measurable subset of a Euclidean space \mathbb{R}^ℓ with $\lambda^\ell(B) > 0$ where λ^ℓ is Lebesgue measure on \mathbb{R}^ℓ . Then a measurable subset A of B is considered generic relative to B , if $\lambda^\ell(B \setminus A) = 0$. This definition generalizes to other spaces B endowed with a canonical measure. However, the space B may lack a canonical measure or, at least, it may be difficult to agree on one. The concept of “prevalence” in Anderson and Zame (1997) extends the measure-theoretical notion of genericity to infinite dimensional parameter sets B that are convex subsets of a topological vector space where an analogue of Lebesgue measure is unavailable.

An alternative is a topological notion of genericity provided that B is a topological space. In the topological context, $A \subset B$ is called generic, if A is open and dense relative to B . The question of what constitutes a natural topology for the economic model is, in part, a pragmatic one. See Hildenbrand (1974; p. 96) and Grandmont et al. (1974).

Since, in our case, B is a Euclidean space, both concepts of genericity are applicable. But as a rule, they do not coincide.⁶ In order to facilitate comparisons with Wilson (1971), Harsanyi (1973), and others, we therefore opt for the measure-theoretic notion with Lebesgue measure as the canonical measure. Whether our main result, Theorem 1 would go through

⁵The alternating move Prisoner's Dilemma game does not.

⁶For example, the set of rational numbers in $B = [0, 1]$ has Lebesgue measure zero. Therefore, the set A of irrational numbers in B is a measure-theoretically generic subset of B . This implies that A is dense in B . Yet, A fails to be open in B and, hence, is not a topologically generic subset of B *strictu sensu*. However, A is a residual subset of B . On the other hand, an open and dense subset of B need not be measure-theoretically generic. Such an example can be obtained through a Cantor-like construction: Cf. Example 4.5 in Gilles, et al. (1998).

as well when the topological notion were used, is an open question. For more on genericity notions, see Mas-Colell (1985).

3 The General Model

3.1 Stochastic Games

Here we define the general class of stochastic games with finite states. Let $N = \{1, \dots, n\}$ denote a finite set of players. Let $\Omega = \{\omega^1, \dots, \omega^K\}$ denote a finite set of states. In each state ω , denote a finite set of actions by $S_{i\omega}$ for each player i . Without loss of generality, we will assume that the sets $S_{i\omega}$ which index state-contingent actions are mutually disjoint so that an action s_i (sometimes denoted $s_{i\omega}$ for clarity) is known to come from a particular $S_{i\omega}$.

Note that standard formulations of stochastic games typically specify state-invariant action sets. Ordinarily, this entails no loss of generality since the actions sets can always be enlarged without changing the strategic properties of the game. However, our concern is with genericity, and so the addition of redundant actions is problematic as it changes the maximal dimension of the game. Hence, our formulation requires state-dependence.

The set of action profiles in state ω is given by $S_\omega \equiv S_{1\omega} \times \dots \times S_{n\omega}$, and $S_{-i\omega} = \times_{j \neq i} S_{j\omega}$. Let $m_{i\omega} = |S_{i\omega}|$ and $m_\omega = |S_\omega| = \prod_{i=1}^n m_{i\omega}$ denote the respective cardinalities. To simplify notation, we will write $m^j \equiv m_{\omega^j}$ for $j = 1, \dots, K$. Finally, set $S = \cup_\omega S_\omega$ and $m = |S| = \sum_{\omega \in \Omega} m_\omega$. Action profiles are denoted by s , or when necessary by s_ω .

Let $U_i(s_\omega)$ denote i 's payoff of action profile $s_\omega \in S_\omega$ in state ω . Player i 's payoffs can be summarized by a $m \times 1$ state dependent column vector U_i given by

$$U_i = [U_i(s)]_{s \in S}^T$$

where $[\cdot]^T$ denote the transpose of the vector. Let $U = (U_i) \in \mathbb{R}^{nm}$ denote the payoff vector of all the players.

The state evolves over time according to a transition function $Q : S \rightarrow \Delta(-)$. We write $Q(\omega' | s_\omega)$ to denote the probability that the game enters state ω' given that the current action profile is $s_\omega \in S_\omega$ in state ω .

Time is assumed discrete, $t = 0, 1, 2, \dots$, and each player i is assumed to use a discount factor $\delta_i \in (0, 1)$. There is an initial distribution π of states at $t = 0$ given by the $1 \times K$

vector

$$\pi = [\pi(\omega^1), \dots, \pi(\omega^K)]$$

In particular, let π_ω denote the Dirac vector that assigns all mass to state ω . A *stochastic game* Γ is summarized by the tuple $\langle N, -, S, U, Q, \delta_1, \dots, \delta_n, \pi \rangle$.

Notice that a standard repeated game has a trivial stochastic game representation in which there is a single state, $- = \{\omega\}$.⁷ A simple example of a nontrivial stochastic game is a resource extraction game such as the one described in Dutta (1996). There are two firms: $n = 2$, and the state ω describes the current amount of a natural resource remaining (e.g., the amount of oil left in the ground). The extraction rate for firm i is given by s_i which is assumed to be chosen from a finite grid in $[0, 1]$. The transition is deterministic: $Q(\omega' | s_\omega) = 1$ whenever $\omega' = [1 - \theta(s_1 + s_2)]\omega$ where $\theta(\cdot)$ is a function of the two firms' extraction rates and $(s_1, s_2) = s_\omega \in S_\omega$. Finally, i 's payoff is his profit given by

$$U_i(s_\omega) = [p((s_1 + s_2)\omega) - c(s_i)]s_i\omega$$

where $p(\cdot)$ is the inverse demand, $c(\cdot)$ is the firm's marginal cost as a function of the extraction rate, and $(s_1, s_2) = s_\omega \in S_\omega$.

Since our concern is in how the equilibrium set depends on payoffs, we use the notation $\Gamma(U)$ to account for the explicit dependence on payoffs. We refer to the *stochastic game form*, $\Gamma(\cdot)$, as the stochastic game structure without having specified payoffs.

Let $s(t)$ denote the profile of actions in period t . A standard notation denotes the complete history of play at time $t + 1$ by $h^{t+1} = (s(0), s(1), \dots, s(t))$. A *behavior strategy* for player i maps each history h and each state ω into a mixed strategy in $\Delta(S_{i\omega})$. A *perfect equilibrium* is a profile of behavior strategies such that each individual's behavior strategy maximizes the expected discounted payoff $E_t \left[\sum_{\tau=t}^{\infty} (1 - \delta_i) \delta_i^{\tau-t} U_i(s(\tau)) \mid h^t, \omega \right]$ after every history h^t and in every state ω which can be reached from h^t via Q .

A result of Dutta (1995) proves a Folk Theorem for the general class of stochastic games. He shows that any payoff which dominates a certain long-run minimax can be arbitrarily approximated by perfect equilibrium payoffs for sufficiently patient players if the individually rational set of payoffs does not vary across histories, and if the game satisfies a full dimensionality requirement.⁸

⁷Equivalently, there may be many states, but Q places all mass on a single state regardless of the conditioning state and profile.

⁸In stochastic games the relevant minmax must be taken over behavior strategies rather than stage game strategies. Dutta defines a limiting minmax as $\delta_i \rightarrow 1$ in order to obtain a reference point which is invariant to discount factor. His result requires that the individually rational payoff set (relative to this minmax) is invariant across histories and that its convex hull have full dimension with respect to the number of players.

The application of Dutta's result here generically resolves the multiplicity issue in stochastic games when all perfect equilibria are considered. However, we will restrict our attention to the special class of strategies known as *Markovian strategies*. These are strategies that depend only on payoff relevant information.

3.2 Markovian Behavior

In the stochastic game the payoff relevant part of history is the current state. A *Markovian* strategy for player i is given by $f_i \in \times_{\omega} \Delta(S_{i\omega})$. We write $f_i(s_i|\omega)$ to denote the probability that i assigns to $s_i \in S_{i\omega}$ when the current state is ω . The players' choices of Markovian strategies induce a joint Markov distribution f defined by

$$f(s|\omega) = \prod_{i=1}^n f_i(s_i|\omega)$$

Let F denote the set of all joint Markov distributions. Given a joint Markov strategy profile $f \in F$ the dynamic payoff $V_i(f|\omega)$ to player i in state ω is expressed recursively by

$$V_i(f|\omega) = \sum_{s \in S_{\omega}} f(s|\omega) \left[(1 - \delta_i) U_i(s) + \delta_i \sum_{\omega'} V_i(f|\omega') Q(\omega'|s) \right]. \quad (1)$$

A *Markov Perfect equilibrium (MPE)* of stochastic game Γ is Perfect equilibrium in Markov strategies. Existence of MPE in finite state stochastic games is well established (see, for example, Friedman (1986, Theorem 4.2)). It is also well known that one best response to Markovian strategies $f_{-i} = (f_j)_{j \neq i}$ is a Markovian strategy. Therefore, f is an MPE whenever

$$V_i(f|\omega) \geq V_i(f_{-i}, \hat{f}_i|\omega), \quad \forall i, \forall \hat{f}_i, \forall \omega$$

4 Generic Finiteness of MPE

4.1 The Main Result

Fix a stochastic game form $\Gamma(\cdot)$, and for each payoff vector $U \in \mathbb{R}^{nm}$, let $E(U)$ denote the set of MPE of $\Gamma(U)$. The main result of the paper is:

Theorem: *Given $\Gamma(\cdot)$, there exists a set $\mathcal{U} \subseteq \mathbb{R}^{nm}$ such that*

1. *the set \mathcal{U} has full Lebesgue measure in \mathbb{R}^{nm} and*
2. *for all $U \in \mathcal{U}$, the set $E(U)$ of MPE of $\Gamma(U)$ is a finite set.*

We remark, first, that the conclusion of the theorem does not depend on the discount factors. The set of MPE is finite regardless of how patient are the players. Naturally, the players' patience does determine which finite set of payoffs can be supported. Maskin and Tirole (1997) prove a genericity result for MPE of finite horizon stochastic games. This is equivalent to setting $\delta_i = 0$ after some time t . Their notion of genericity applies to the fully dimensional state dependent payoffs. Hence, their result is a finite horizon analogue of ours. For this reason, the backward induction techniques used in their proof cannot be applied here.

Note, second, that in the standard repeated game, the set of MPE coincide with repetitions of Nash equilibria of the stage game. Hence, the result implies the Harsanyi-Rosenmüller-Wilson extension to MPE of repeated games.

The proof is organized into three key steps. The first establishes a genericity result for the closure of the set of *completely mixed* MPE. An MPE f is completely mixed if the joint Markov profile satisfies, for each ω and each $s \in S_\omega$, $f(s|\omega) > 0$. Let $E^\circ(U)$ denote the set of completely mixed MPE of $\Gamma(U)$. Let $\overline{E^\circ(U)}$ denote the closure of $E^\circ(U)$. The first step then shows:

Proposition: *Given $\Gamma(\cdot)$, there exists a set \mathcal{U} with full measure in \mathbb{R}^{nm} such that for each $U \in \mathcal{U}$, the set $\overline{E^\circ(U)}$ consists of isolated points in F , the set of joint Markov strategies.*

The second step replaces “isolated points” by “finitely many points”, utilizing a standard compactness argument. The final step yields the Theorem as a corollary. The intuition behind it is that MPE that are not completely mixed are, in fact, completely mixed on a smaller support. Hence, we examine the reduced stochastic game with restricted support, and note that the previous results apply. Since there are only finitely many supports, generic finiteness holds for all MPE. The next two subsections introduce relevant notation and outline the logic of the proof.

4.2 Some Notation

For our purposes, it will prove more convenient to express payoffs in terms of sums and products of matrices. First, for any joint Markov distribution f let \mathcal{P}_f denote the $K \times m$

matrix induced by f :

$$\mathcal{P}_f = \begin{bmatrix} f(s_{\omega^1}^1|\omega^1) & \dots & f(s_{\omega^1}^{m_1}|\omega^1) & 0 & \dots & \dots & 0 \\ 0 & \dots & 0 & f(s_{\omega^2}^1|\omega^2) & \dots & f(s_{\omega^2}^{m_2}|\omega^2) & \dots & \vdots \\ \vdots & & & & & \vdots & \ddots & 0 \\ 0 & \dots & & \dots & & & f(s_{\omega^K}^1|\omega^K) & \dots & f(s_{\omega^K}^{m_K}|\omega^K) \end{bmatrix}$$

Next, write Q as a $m \times K$ transition matrix for the stochastic game, given by

$$Q = \begin{bmatrix} Q(\omega^1|s_{\omega^1}^1), & Q(\omega^2|s_{\omega^1}^1) & \dots & Q(\omega^K|s_{\omega^1}^1) \\ Q(\omega^1|s_{\omega^1}^2) & \dots & \dots & \\ \vdots & \vdots & & \\ Q(\omega^1|s_{\omega^1}^{m_1}) & \vdots & & \vdots \\ Q(\omega^1|s_{\omega^2}^1) & \vdots & & \\ \vdots & \vdots & & \\ \vdots & \vdots & & \vdots \\ Q(\omega^1|s_{\omega^K}^{m_K}) & \dots & \dots & Q(\omega^K|s_{\omega^K}^{m_K}) \end{bmatrix}$$

Now, the dynamic payoffs in $\Gamma(U)$ can be expressed as linear functions of payoff parameters:

$$\begin{aligned} V_i(f|\omega) &= \pi_\omega \cdot \left[\sum_{t=0}^{\infty} (1 - \delta_i) \delta_i^t (\mathcal{P}_f \cdot Q)^t \cdot \mathcal{P}_f \right] \cdot U_i \\ &\equiv \pi_\omega \cdot \mathcal{A}_f \cdot U_i \end{aligned} \tag{2}$$

Clearly, the right-hand side of (2) is smooth in the pair (f, U_i) . But notice that for player i and state ω , the probabilities $f_i(s_i|\omega), s_i \in S_{i\omega}$, have to satisfy $\sum_{s_i} f_i(s_i|\omega) = 1$, hence cannot vary independently. For the case of completely mixed MPE — which will be assumed in the first step — one gets around this problem by fixing some action $d_{i\omega} \in S_{i\omega}$ for player i . Let $D_i = \{d_{i\omega} : \omega \in -\}$ be denoting the set of “normalizing actions” for player i . By varying $\omega \in -$ and $i = 1, \dots, n$, we pin down Kn normalizing equations satisfying

$$f_i(d_{i\omega}|\omega) = 1 - \sum_{s_i \in S_{i\omega} \setminus \{d_{i\omega}\}} f_i(s_i|\omega), \quad \omega \in -, \quad i = 1, \dots, n. \tag{3}$$

Finally, for $s_i \notin D_i$, the $K \times m$ matrix

$$\mathcal{B}_f(s_i|\omega) = \frac{\partial}{\partial f_i(s_i|\omega)} [\mathcal{A}_f]$$

is especially relevant to the proof.

4.3 An Outline of the Argument

Now we are ready to draw a more detailed road map for the proof. The first step, showing the Proposition, is the critical one. To understand its logic, it is useful to compare it to Harsanyi's, by now standard, result for normal form games. He obtains a system of equations by examining first order necessary conditions for Nash equilibrium. In normal form games, strategies enter linearly in payoffs. Harsanyi exploits this linearity to derive an explicit mapping from equilibrium strategies to payoffs. Sard's Theorem asserts that the regular values of this map have full measure. The linearity of the system is again used to show that the corresponding regular strategies are locally isolated in the simplex.

Like Harsanyi, we also proceed from a system of equations derived from first order necessary conditions, $D_f[V_i] = 0$, for Markov Perfection. We then would have to show full rank of the Jacobian matrix $D_U[D_f[V_i]]$. It turns out that the first order necessary conditions translate into a system of $\bar{m} - Kn$ equations

$$[\pi_\omega \cdot \mathcal{B}_f(s_i|\omega) \cdot U_i] = 0 \quad (4)$$

where $i = 1, \dots, n$, $\omega \in \Omega$, $s_i \notin D_i$, $\bar{m} \equiv \sum_i \bar{m}_i \equiv \sum_i \sum_\omega m_{i\omega}$ (this step is Lemma 1). To economize on notation, we rewrite the system of equations as a single vector equation $M(f, U) = 0$ that is satisfied for any pair (f, U) with $f \in \overline{E^\circ(U)}$ (this step is Lemma 2 (i)). The main task at hand is to show that $\partial M/\partial U$ has full rank (Lemma 2 (ii)). However, here a difficulty arises which prevents us from simply adapting Harsanyi's result. Unlike in normal form games, our system is not linear. Rather, players' dynamic payoffs in stochastic games are shown to be rational functions of Markovian strategies. This necessitates a lengthy and tedious proof of the full rank condition on the system $M(\cdot, \cdot)$ perturbed by U .

Since, roughly, M maps $F \times \mathbb{R}^{nm}$ to $\mathbb{R}^{\bar{m}-Kn}$, we can apply the Implicit Function Theorem to derive (locally) an implicit, smooth map Ψ from strategies, i.e. F to an $\bar{m} - Kn$ -dimensional subspace of \mathbb{R}^{nm} so that points on the graph of this map satisfy the first order necessary conditions for MPE (Lemma 3).

From there, the logic (if not the details) is more standard. Sard's and Fubini's Theorems are used to show that regular values of Ψ have full measure in \mathbb{R}^{nm} and that the corresponding regular points are locally isolated (Lemmas 4 and 5). Throughout most of the proof, until and including Lemma 5, local analysis is performed. A subsequent countability argument yields a global result as asserted in the Proposition. As pointed out earlier, two more steps remain to complete the proof of the Theorem.

An alternative route to genericity results has been taken by Balasko (1988) who performs global analysis from the very beginning. While more elegant, it would not spare us the

tedious proof of the full rank assertion of Lemma 2 below. A second alternative route to genericity results has been taken recently by Blume and Zame (1994) and Govindan and McLemman (1998a,b) who rely on the theory of semi-algebraic sets. A semi-algebraic set is a set characterized by a finite system of polynomial inequalities. Though it is not obvious, the implicit function Ψ in our proof can apparently be shown to be semi-algebraic.⁹ It turns out that semi-algebraic systems have appealing generic properties (the Blume and Zame paper is a good reference). It appears to us that these properties could be used to establish that the inverse image under Ψ of regular values of Ψ is a zero dimensional set in the space of Markov profiles f . This would make the alternative route attractive since it simplifies the latter part of our proof (though not the cumbersome Lemma 2) and possibly yields a stronger genericity result.¹⁰ By preferring the more familiar perturbation methods of Debreu (1970) and Harsanyi (1973), we make the local arguments explicit rather than embed them in the semi-algebraic results.

4.4 A Proof

Let $\Gamma(\cdot)$ be as hypothesized. We are going to prove the Proposition first. Choose a normalization $D = (D_1, \dots, D_n)$. Notice that F is a compact and full-dimensional convex subset of the $\bar{m} - Kn$ -dimensional affine subspace of $\mathbb{R}^{\bar{m}}$ given by (3).

Lemma 1 *For any $U \in \mathbb{R}^{nm}$, let $f \in E^\circ(U)$. Then it satisfies each of the $\bar{m} - Kn$ equations described by (4), subject to (3).*

The proofs of this and all subsequent Lemmata are contained in the Appendix.

We now define the following neighborhood of F . Let $\epsilon > 0$ and define the open set F^ϵ such that $f \in F^\epsilon$ implies

- (i) for each i , $-\epsilon < f_i(s_i|\omega)$ for $s_i \in S_{i\omega}$, $\omega \in \cdot$.
- (ii) $\sum_{S_{i\omega}} f_i(s_i|\omega) < 1 + \epsilon$.

Clearly F^ϵ contains the compact set F where the topological notions are relative to the affine subspace given by (3). For sufficiently small ϵ , all the subsequent differential geometry and

⁹A function is semi-algebraic if its graph is a semi-algebraic set.

¹⁰One of the referees outlines an argument which uses semi-algebraic methods on the Bellman equation (1). If successful, this might lead to a still shorter proof.

topology arguments are valid on F^ϵ . By extending the analysis to F^ϵ for the time being, we allow the values $f_i(s_i|\omega)$ and $f_i(d_{i\omega}|\omega)$ to become zero or slightly negative. We do, however, maintain the restrictions (3). Replacing ϵ by 0 in (i) and (ii) yields the convenient notation F° for the interior of F which contains the completely mixed MPE.

From now on, we work with the mapping M associated with the left-hand side of (4), i.e. $M : F^\epsilon \times \mathbb{R}^{nm} \rightarrow \mathbb{R}^{\bar{m}-Kn}$ where

$$M_{i s_{i\omega}}(f, U) = \pi_\omega \cdot \mathcal{B}_f(s_{i\omega}|\omega) \cdot U_i. \quad (5)$$

Lemma 2 *Let $U \in \mathbb{R}^{nm}$. Then:*

- (i) $M(f, U) = 0$ for all $f \in \overline{E^\circ(U)}$.
- (ii) $\partial M/\partial U$ has full rank $\bar{m} - Kn$ at (f, U) for all $f \in F^\epsilon$.

Lemmata 1 and 2 are used to establish the next three Lemmata which collectively lay the groundwork for local properties of profiles in $\overline{E^\circ(U)}$. Call $C = (C^1, C^2)$ a *suitable partition* of $\{1, \dots, m\}$, if $C^2 \cup C^1 = \{1, \dots, m\}$, and $C^2 \cap C^1 = \emptyset$, with $|C^2| = \bar{m} - Kn$. Given a suitable partition C , we decompose the vectors $U \in \mathbb{R}^{nm}$ accordingly: $U = (U_{C^1}, U_{C^2})$.

Lemma 3 *Let $(f^0, U^0) \in F^\epsilon \times \mathbb{R}^{nm}$ satisfy $M(f^0, U^0) = 0$. Then there exist a suitable partition $C = (C^1, C^2)$ and*

- (i) *open neighborhoods \mathcal{N}_F , \mathcal{N}_{C^1} and \mathcal{N}_{C^2} of f^0 , $U_{C^1}^0$ and $U_{C^2}^0$, respectively, and*
- (ii) *a locally smooth mapping $\Psi : \mathcal{N}_F \times \mathcal{N}_{C^1} \rightarrow \mathcal{N}_{C^2}$*

such that if $U_{C^2} = \Psi(f, U_{C^1})$ then U_{C^2} is the unique subvector in \mathcal{N}_{C^2} which solves the equation $M(f, U_{C^1}, U_{C^2}) = 0$.

Now let again $(f^0, U^0) \in F^\epsilon \times \mathbb{R}^{nm}$ with $M(f^0, U^0) = 0$. Fix the neighborhoods \mathcal{N}_F , \mathcal{N}_{C^2} and \mathcal{N}_{C^1} and the mapping Ψ obtained in Lemma 3. Call $(f, U_{C^1}) \in \mathcal{N}_F \times \mathcal{N}_{C^1}$ an *f-regular point* of Ψ , if $\partial \Psi/\partial f$ has full rank $\bar{m} - Kn$ at (f, U_{C^1}) . Call $(f, U_{C^1}) \in G$ an *f-critical point* of Ψ otherwise. Call $U_{C^2} \in \mathcal{N}_{C^2}$ an *f-critical value* of Ψ , if $U_{C^2} = \Psi(f; U_{C^1})$ with $(f; U_{C^1})$ an *f-critical point* of Ψ .

Lemma 4 *There exists a correspondence $Z : \mathcal{N}_{C^1} \dashrightarrow \mathcal{N}_{C^2}$ such that for each $U_{C^1} \in \mathcal{N}_{C^1}$,*

- (i) *the set $Z(U_{C^1})$ has Lebesgue measure zero in \mathcal{N}_{C^2} and*
- (ii) *the set of f -critical values of $\Psi(\cdot, U_{C^1})$ coincides with $Z(U_{C^1})$.*

Clearly, the Lemma implies that the set of f -critical values of $\Psi(\cdot, U_{C^1})$ has Lebesgue measure zero in \mathcal{N}_{C^2} . Now fix $\mathcal{N}_F, \mathcal{N}_{C^2}, \mathcal{N}_{C^1}$ with $(f^0, U^0) \in \mathcal{N}_F \times \mathcal{N}_{C^2} \times \mathcal{N}_{C^1}$ as before with $Z(\cdot)$ as in Lemma 4.

Lemma 5 *Let $\mathcal{Z} = \{(U_{C^1}, U_{C^2}) \in \mathcal{N}_{C^1} \times \mathcal{N}_{C^2} \mid U_{C^2} \in Z(U_{C^1})\}$ be the graph of Z .*

- (i) *\mathcal{Z} is measurable and has Lebesgue measure zero in $\mathcal{N}_{C^1} \times \mathcal{N}_{C^2}$.*
- (ii) *If $U = (U_{C^1}, U_{C^2}) \in \mathcal{N}_{C^1} \times \mathcal{N}_{C^2}$ and $U \notin \mathcal{Z}$, then the elements of $\mathcal{N}_F(U) \equiv \{f \in \mathcal{N}_F : M(f, U) = 0\}$ are isolated points.*

We now utilize the results of the previous three lemmata to establish the Proposition.

Proof of Proposition We assume without loss of generality that \mathcal{N}_F is an open box with rational valued end-points. Similarly, we assume $\mathcal{N}_C \equiv \mathcal{N}_{C^1} \times \mathcal{N}_{C^2}$ to be an open box in \mathbb{R}^{nm} , again with rational valued end-points. Consequently, application of Lemmata 3-5 to all qualifying pairs (f^0, U^0) determines a sequence of triples $(\mathcal{N}_F^\tau, \mathcal{N}_C^\tau, \mathcal{Z}^\tau)$, $\tau = 1, 2, \dots$, such that

- (a) for each pair $(f^0, U^0) \in F^\epsilon \times \mathbb{R}^{nm}$ with $M(f^0, U^0) = 0$, there exists some τ with $(f^0, U^0) \in \mathcal{N}_F^\tau \times \mathcal{N}_C^\tau$;
- (b) for each τ , \mathcal{N}_F^τ is a (relatively) open subset of F^ϵ ;
- (c) for each τ , \mathcal{N}_C^τ is an open subset of \mathbb{R}^{nm} ;
- (d) for each τ , \mathcal{Z}^τ is a measure zero subset of \mathcal{N}_C^τ ;
- (e) for each τ and $U \in \mathcal{N}_C^\tau \setminus \mathcal{Z}^\tau$, the elements of $\mathcal{N}_F^\tau(U)$ are isolated points.

Next define $\mathcal{Z}_D^\infty = \bigcup_\tau \mathcal{Z}^\tau$ where the subscript D indicates that the set depends on the normalization $D = (D_1, \dots, D_n)$. Observe that \mathcal{Z}_D^∞ has Lebesgue measure zero in \mathbb{R}^{nm} . There are finitely many conceivable normalizations. Let \mathcal{Z}^∞ be the union of the finitely many sets \mathcal{Z}_D^∞ , i.e., $\mathcal{Z}^\infty = \bigcup_D \mathcal{Z}_D^\infty$. Then \mathcal{Z}^∞ has also Lebesgue measure zero in \mathbb{R}^{nm} . Therefore, the set $\mathcal{U} \equiv \mathbb{R}^{nm} \setminus \mathcal{Z}^\infty$ has full measure in \mathbb{R}^{nm} .

Now fix $U \in \mathcal{U}$. By Lemma 2(i), $\overline{E^\circ(U)} \subseteq \bigcup_\tau \mathcal{N}_F^\tau(U)$, and so by Lemma 5 the points in $\overline{E^\circ(U)}$ are isolated points in F^ϵ . \square

Proof of Theorem Utilizing the same full measure set \mathcal{U} constructed in the proof of the Proposition, we claim that $\overline{E^\circ(U)}$ is finite for all $U \in \mathcal{U}$. Fix any $U \in \mathcal{U}$. Since $\overline{E^\circ(U)}$ is a closed subset of the compact set F , it is also compact. It can therefore be covered by a family of open, non-empty, pairwise disjoint sets $N_f, f \in \overline{E^\circ(U)}$ in which $N_f \cap \overline{E^\circ(U)} = \{f\}$. By compactness, there exists a finite sub-covering. But since each N_f contains no points in $\overline{E^\circ(U)}$ other than f , this is only possible if $\overline{E^\circ(U)}$ itself is finite. Since the choice of $U \in \mathcal{U}$ is arbitrary, the claim is shown.

It remains to deal with MPE that are not completely mixed. Let $T \subseteq S$ denote an arbitrary support for a joint Markov profile. That is, T is of the form $T = \bigcup_\omega \times_i T_{i\omega}$ where $T_{i\omega} \subseteq S_{i\omega}$, each i and ω . Let F^T denote the compact subsimplex of F that contains all Markov profiles with support T :

$$F^T = \{f \in F : s \in S_\omega \setminus T \Rightarrow f(s|\omega) = 0\}$$

Let \mathbb{R}^T denote the lower dimensional Euclidean space that obtains when S is replaced by T . Then $\mathbb{R}^{nm} = \mathbb{R}^T \times \mathbb{R}^{T^c}$. For $U_T \in \mathbb{R}^T$, consider the restricted stochastic game $\Gamma^T(U_T) = \langle \cdot, T, U_T, Q_T, \pi \rangle$. Let $E^\circ(U_T) \subset F^T$ denote the set of joint profiles in F^T whose restrictions to T are completely mixed MPE of the reduced game $\Gamma^T(U_T)$. By the Proposition, there exists a generic subset \mathcal{U}_T of \mathbb{R}^T such that for all $U_T \in \mathcal{U}_T$, $\overline{E^\circ(U_T)}$ consists of locally isolated points in F^T , and so by the above compactness argument, it is a finite set. By Fubini's Theorem, $\mathcal{U}_T \times \mathbb{R}^{T^c}$ is a generic subset of \mathbb{R}^{mn} .

Now for $U \in \mathcal{U}_T \times \mathbb{R}^{T^c}$, let U_T denote its restriction to T . Then $f \in E(U)$ has support T if and only if $f \in E^\circ(U_T)$. Consequently, $\Gamma(U)$ has finitely many MPE with support T . However, there are only finitely many supports. Hence, the assertion of the Theorem holds. \square

5 Lower Dimensional Games

In light of our result several issues should be addressed. First, while the theorem covers generic “full dimensional” stochastic games in \mathbb{R}^{nm} , there are a number of interesting stochastic games which have been studied in the literature that are not covered by the result. A case in point is the Markov analysis of two-player alternating move games examined in a series of papers by Maskin and Tirole (1987, 1988a, 1988b). For instance, in the (1988a) paper, each firm sets capacity given the temporarily fixed capacity of rivals. See also Lagunoff and Matsui (1997).

The alternating move game has a straightforward stochastic game representation: fix a stage game with actions sets S_1 and S_2 , and with $u_i(s)$ denoting i 's stage game payoff when profile $s \in S_1 \times S_2$ is chosen. At the beginning of odd periods, player 1 has a chance to revise his action, whereas at the beginning of even periods player 2 has a chance to revise her own action. Formally, let $S = S_1 \cup S_2$. Further let $S_{i\omega} = \{\omega\}$, if $\omega \in S_i$, and $S_{i\omega} = S_i$, if $\omega \notin S_i$. Q is given by $Q(\omega'|s_\omega) = 1$ in case $\omega \in S_i$, $\omega' \in S_j$, $s_{j\omega} = \omega'$, $i \neq j$; and by $Q(\omega'|s_\omega) = 0$ otherwise. The state dependent payoffs assume the form

$$U_i(s_\omega) = \begin{cases} u_i(\omega, s_{j\omega}) & \text{if } \omega \in S_i; \\ u_i(s_{i\omega}, \omega) & \text{if } \omega \notin S_i. \end{cases}$$

To see why the existing result does not apply in these games, consider an alternating move game derived from the stage game in Figure 1 below.

II

		s_2	s'_2
I	s_1	5, 3	20, 2
	s'_1	10, 1	0, 4
	s''_1	11, 6	9, 18

Figure 1

Letting arrows denote the transition laws, Figure 2 exhibits the stochastic game:

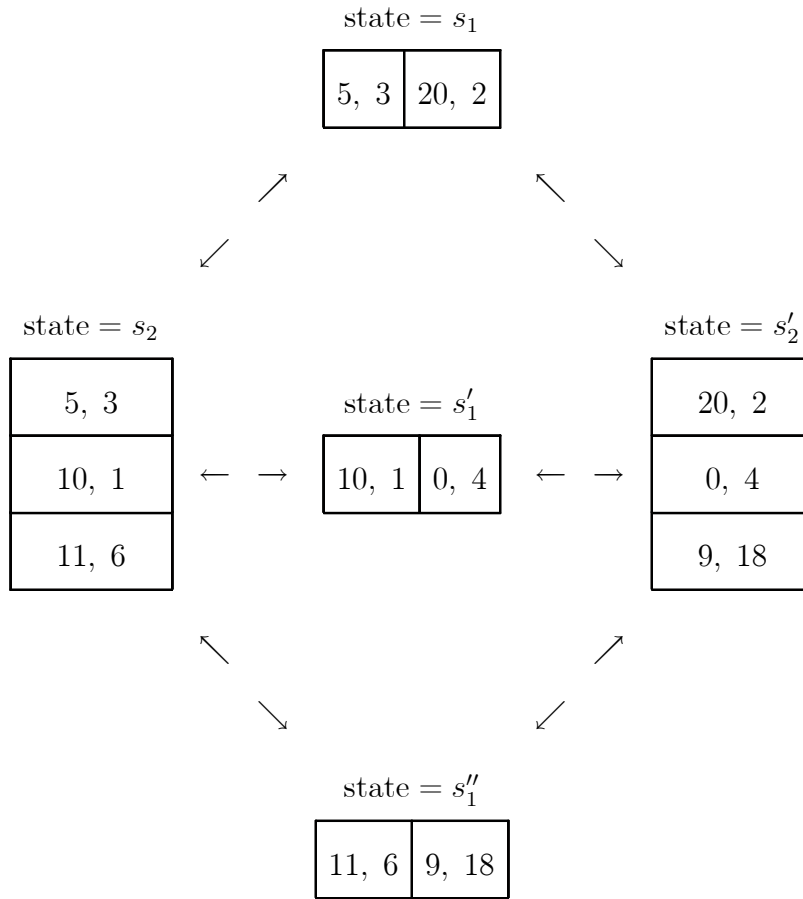


Figure 2

In this example, the state dependent payoff vectors reside in \mathbb{R}^{24} . But in \mathbb{R}^{24} , the representation of repeated alternating moves in Figure 2 is clearly *not* a “generic” stochastic game; the alternating move game admits not more than 12 degrees of freedom — the dimension of the original stage game. Clearly, in the alternating move game, “genericity” should refer to the lower dimensional payoffs of the original stage game. In a related paper, Haller and Lagunoff (1997), we show that the set of MPE of the 2-player, alternating move game is indeed generically finite with respect to stage game payoffs. For other types of games and equilibrium concepts, some lower dimensional genericity results also exist. Examples are Govindan and McLennan (1998b) for extensive game forms, and Park (1997) for signalling games.

6 Appendix

Proof of Lemma 1 Let $U \in \mathbb{R}^{nm}$ and $f \in E^\circ(U)$. Then for each i , each $\omega \in -$, and $s_i \in S_{i\omega}$,

$$\left. \frac{\partial V_i(f|\omega)}{\partial f_i(s_i|\omega)} \right|_{\sum_{t_i} f_i(t_i|\omega)=1} = 0.$$

Since the dynamic payoff functions V_i assume the form (2) and the normalization D is chosen, this system of equations translates into the $\bar{m} - Kn$ equations

$$f_i(s_i|\omega) f_i(d_{i\omega}|\omega) [\pi_\omega \cdot \mathcal{B}_f(s_i|\omega) \cdot U_i] = 0$$

where $s_i \notin D_i$ and (3) holds. Since $f \in E^\circ(U)$, we can drop the multipliers $f_i(s_i|\omega) f_i(d_{i\omega}|\omega)$ and obtain (4). \square

Proof of Lemma 2 Let $f^k \rightarrow f$ satisfy $f^k \in E^\circ(U)$ and hence by Lemma 1 and (5), $M(f^k, U) = 0$ for each k . Since M is smooth, $M(f, U) = 0$ which shows assertion (i).

Fix a pair $(f, U) \in F^\epsilon$. We begin by analyzing the derivative of $M(f, U)$ with respect to U , given by the Jacobian $D_U M(f, U)$. This Jacobian is a $(\bar{m} - nK) \times m$ matrix which has the form:

$$D_U M(f, U) = \begin{bmatrix} J_U^{11}(f) & 0 & \cdots & 0 \\ 0 & J_U^{22}(f) & & \vdots \\ \vdots & & \ddots & \\ 0 & \cdots & \cdots & J_U^{nn}(f) \end{bmatrix}$$

Figure 4

Here, $J_U^{ii}(f)$ denotes the $(\bar{m}_i - K) \times m$ block of partials

$$\frac{\partial M_{i s_{i\omega}}(f, U)}{\partial U_i(s')} = \frac{\partial [\pi_\omega \cdot \mathcal{B}_f(s_{i\omega}|\omega) \cdot U_i]}{\partial U_i(s')}, \quad s_{i\omega} \in S_{i\omega}, \quad i = 1, \dots, n, \quad s' \in S \quad (6)$$

There are zeroes outside of the $J_U^{ii}(f)$ blocks since the partials vanish wherever $j \neq i$. Given the linearity of payoffs in U_i each entry in $J_U^{ii}(f)$ assumes the form

$$\frac{\partial M_{i s_{i\omega}}(f, U)}{\partial U_i(s')} = \pi_\omega \cdot \mathcal{B}_f(s_{i\omega}|\omega) \cdot 1_{s'} \quad (7)$$

where $s_{i\omega} \in S_{i\omega}$, $\omega \in -$ and $1_{s'}$ is the $m \times 1$ Dirac vector with a 1 in the s' component and 0's elsewhere. To be clear, the rows of $J_U^i(f)$ vary across actions $s_{i\omega}$ while the columns vary across action profiles $s \in S$. Note that, due to the linearity of dynamic payoffs in the vector U , vector U does not appear in (7). Hence, $J_U^i(f)$ is expressed as a function of f only. We first show that for each i , $J_U^i(f)$ has full row rank $\bar{m}_i - K$.

Observe that

$$\begin{aligned}
\mathcal{B}_f(s_{i\omega}|\omega) &= \frac{\partial}{\partial f_i(s_i|\omega)} [\mathcal{A}_f] \\
&= \frac{\partial}{\partial f_i(s_i|\omega)} \left[(1 - \delta_i) [I - \delta_i \mathcal{P}_f \cdot Q]^{-1} \cdot \mathcal{P}_f \cdot \right] \\
&= (1 - \delta_i) \left([I - \delta_i \mathcal{P}_f \cdot Q]^{-1} \cdot \left[\frac{\partial \mathcal{P}_f}{\partial f_i(s_i|\omega)} \right] \cdot \right. \\
&\quad \left. - [I - \delta_i \mathcal{P}_f \cdot Q]^{-1} \cdot \right. \\
&\quad \left. \cdot \left[-\delta_i \frac{\partial \mathcal{P}_f}{\partial f_i(s_i|\omega)} \cdot Q \right] \cdot [I - \delta_i \mathcal{P}_f \cdot Q]^{-1} \cdot \mathcal{P}_f \cdot \right) \\
&= (1 - \delta_i) [I - \delta_i \mathcal{P}_f \cdot Q]^{-1} \cdot \left[\frac{\partial \mathcal{P}_f}{\partial f_i(s_i|\omega)} \right] \\
&\quad \cdot \left[\hat{I} + \delta_i Q \cdot [I - \delta_i \mathcal{P}_f \cdot Q]^{-1} \cdot \mathcal{P}_f \right],
\end{aligned} \tag{8}$$

where \hat{I} is the $m \times m$ identity matrix,¹¹ and $\frac{\partial \mathcal{P}_f}{\partial f_i(s_i|\omega)}$ denotes the entry by entry derivative of \mathcal{P}_f with respect to $f_i(s_i|\omega)$. Note that the inverse $[I - \delta_i \mathcal{P}_f \cdot Q]^{-1}$ exists because $\mathcal{P}_f \cdot Q$ is a stochastic matrix and so $I - \delta_i \mathcal{P}_f \cdot Q$ is a matrix with a strictly dominant diagonal. Combining our knowledge of the entries of $J_U^i(f)$ in (7) with equation (8), it follows that $J_U^i(f)$ equals

$$(1 - \delta_i) \cdot \begin{bmatrix} \pi_{\omega^1} \cdot [I - \delta_i \mathcal{P}_f \cdot Q]^{-1} \cdot \left[\frac{\partial \mathcal{P}_f}{\partial f_i(s_{i\omega^1}|\omega^1)} \right] \\ \vdots \\ \pi_{\omega^K} \cdot [I - \delta_i \mathcal{P}_f \cdot Q]^{-1} \cdot \left[\frac{\partial \mathcal{P}_f}{\partial f_i \left(\begin{smallmatrix} m^K \\ s_{i\omega^K} \end{smallmatrix} \middle| \omega^K \right)} \right] \end{bmatrix} \cdot \left[\hat{I} + \delta_i Q \cdot [I - \delta_i \mathcal{P}_f \cdot Q]^{-1} \cdot \mathcal{P}_f \right] \cdot \hat{I} \tag{9}$$

To clarify the expression (9), the entries in (7) are post-multiplied by unit vectors $1_{s\omega}$. But these entries form an identity matrix \hat{I} which post-multiplies expression (9).

¹¹The hat on \hat{I} is there to distinguish it from the $K \times K$ identity matrix I .

We first examine the last block in expression (9). Observe:

$$\begin{aligned}
\hat{I} + \delta_i Q [I - \delta_i \mathcal{P}_f \cdot Q]^{-1} \cdot \mathcal{P}_f &= \hat{I} + \delta_i Q \cdot \sum_{t=0}^{\infty} \delta_i^t (\mathcal{P}_f \cdot Q)^t \cdot \mathcal{P}_f \\
&= \hat{I} + \sum_{t=1}^{\infty} \delta_i^t (Q \cdot \mathcal{P}_f)^t \\
&= \sum_{t=0}^{\infty} \delta_i^t (Q \cdot \mathcal{P}_f)^t \\
&= [\hat{I} - \delta_i Q \cdot \mathcal{P}_f]^{-1},
\end{aligned} \tag{10}$$

where, as with $[I - \delta_i \mathcal{P}_f \cdot Q]^{-1}$, the inverse $[\hat{I} - \delta_i Q \cdot \mathcal{P}_f]^{-1}$ exists because $Q \cdot \mathcal{P}_f$ is a stochastic matrix and so $\hat{I} - \delta_i Q \cdot \mathcal{P}_f$ is a matrix with a strictly dominant diagonal. Since $[\hat{I} - \delta_i Q \cdot \mathcal{P}_f]^{-1}$ is invertible it must have full rank of m .

We now turn our attention to the first block in expression (9). The matrix $\frac{\partial \mathcal{P}_f}{\partial f_i(s_i|\omega)}$ is a $K \times m$ matrix placing zeroes in all entries where $f_i(s_i|\omega)$ does not appear, and drops $f_i(s_i|\omega)$ from wherever it appears in the joint Markov profile $f(s|\omega)$. Note that by our earlier normalization, $f_i(s_i|\omega)$ appears in the $d_{i\omega}$ entry since $f_i(d_{i\omega}|\omega) = 1 - \sum_{s'_i \neq d_{i\omega}} f_i(s'_i|\omega)$. Letting f_{-i} denote the joint Markov profile excluding i , and fixing $\omega = \omega^1$ the matrix $\frac{\partial \mathcal{P}_f}{\partial f_i(s_i|\omega^1)}$ has the form:

$$\begin{bmatrix}
0 & \dots & 0 & \dots & f_{-i}(s_{-i}^l|\omega^1) & \dots & 0 & \dots & 0 & \dots & -f_{-i}(s_{-i}^l|\omega^1) & \dots & 0 & \dots & 0 \\
0 & \dots & & & 0 & & & & & & 0 & & & & 0 \\
\vdots & & \vdots & & & & & & & & & & \vdots & & \vdots \\
0 & \dots & & & 0 & & & & & & 0 & & & & 0
\end{bmatrix}$$

where l symbolizes variation over all profiles s with i th component s_i . Clearly, by (3), there is some s_{-i} for which $f_{-i}(s_{-i}|\omega^1) \neq 0$. Hence, the preceding matrix has some nonzero entries in row ω^1 at locations $(s_{i\omega^1}, s_{-i\omega^1}^l)$ and $(d_{i\omega^1}, s_{-i\omega^1}^l)$ and zeroes elsewhere. Let $\Lambda_i(f; s_{i\omega^1})$ denote the $1 \times m$ row vector $\pi_{\omega^1} \cdot [I - \delta_i \mathcal{P}_f \cdot Q]^{-1} \cdot [\frac{\partial \mathcal{P}_f}{\partial f_i(s_i|\omega^1)}]$. Indeed, let $\Lambda_i(f; s_{i\omega})$ denote the corresponding $1 \times m$ row vector for an arbitrary action $s_{i\omega}$. The vector $\Lambda_i(f; s_{i\omega})$ has non-zero entries at certain actions $s_{i\omega}$ where $s_{i\omega}$ or $d_{i\omega}$ is the i th component of s_ω and zero entries elsewhere. Hence the $(\bar{m}_i - K) \times m$ matrix $\Lambda_i(f)$ has $\bar{m}_i - K$ independent columns which means that it has full rank $\bar{m}_i - K$.

Summarizing, we obtain $J_U^{ii}(f)$ as

$$J_U^{ii}(f) = (1 - \delta_i) \Lambda_i(f) \cdot [\hat{I} - \delta_i Q \cdot \mathcal{P}_f]^{-1} \cdot \hat{I}. \tag{11}$$

We now assert that the matrix $J_U^{ii}(f)$ defined in equation (11) has full row rank. To see why, a standard fact about matrix algebra is used for any two matrices A and B whose product is defined: $\text{Rank}(AB) = \text{Rank}(A)$ if B is square and nonsingular. Here we have that $\Lambda_i(f)$ has rank of $\bar{m}_i - K$ while $[\hat{I} - \delta_i Q \cdot \mathcal{P}_f]^{-1} \cdot \hat{I}$ is square and nonsingular. Applying this fact to (11) establishes that $\text{Rank}(J_U^{ii}(f)) = \bar{m}_i - K$. Because of the diagonal box form in Figure 4, $\frac{\partial M}{\partial U}$ has full rank $\bar{m} - nK = \sum_i(\bar{m}_i - K)$ as asserted. This concludes the proof of assertion (ii). \square

Proof of Lemma 3 By Lemma 2, $D_U M(f, U)$ has full rank $\bar{m} - nK$ at (f^0, U^0) . Therefore, we can find a suitable partition (C^1, C^2) such that

$$\frac{\partial M(f, U)}{\partial U_{C^2}} \text{ has full rank } \bar{m} - nK \text{ at } (f^0, U^0).$$

By the Implicit Function Theorem, there exist $\mathcal{N}_F, \mathcal{N}_{C^1}, \mathcal{N}_{C^2}$, and Ψ with the desired properties. \square

Proof of Lemma 4 For each $U_{C^1} \in \mathcal{N}_{C^1}$, the mapping $\Psi(\cdot, U_{C^1})$ from \mathcal{N}_F to \mathcal{N}_{C^2} is smooth in f . Hence, by Sard's theorem, the set of critical values of $\Psi(\cdot, U_{C^1})$ constitutes a set $Z(U_{C^1})$ of Lebesgue measure zero in \mathcal{N}_{C^2} . This defines the set valued mapping $Z(\cdot)$ with the asserted properties. \square

Proof of Lemma 5 The proof of measurability of \mathcal{Z} follows a lead in the proof of Sard's theorem in Milnor (1965), p. 18. First, recall the mapping $\Psi : \mathcal{N}_F \times \mathcal{N}_{C^1} \rightarrow \mathcal{N}_{C^2}$. Let $\mathcal{Y} \subseteq \mathcal{N}_F \times \mathcal{N}_{C^1}$ denote the set of f -critical points in $\mathcal{N}_F \times \mathcal{N}_{C^1}$. Then $\mathcal{Z} = \Phi(\mathcal{Y})$ where the smooth mapping $\Phi : \mathcal{N}_F \times \mathcal{N}_{C^1} \rightarrow \mathcal{N}_{C^1} \times \mathcal{N}_{C^2}$ is defined by $\Phi(f, U_{C^1}) = (U_{C^1}, \Psi(f, U_{C^1}))$.

The set $(\mathcal{N}_F \times \mathcal{N}_{C^1}) \setminus \mathcal{Y}$ is open relative to $\mathcal{N}_F \times \mathcal{N}_{C^1}$ and thus \mathcal{Y} is closed relative to $\mathcal{N}_F \times \mathcal{N}_{C^1}$. Observe $\mathcal{N}_F \times \mathcal{N}_{C^1} \subset \mathbb{R}^{\bar{m}} \times \mathbb{R}^{m-(\bar{m}-nK)} = \mathbb{R}^{m+nK}$. Let \mathbb{Q}^{m+nK} denote the set of points in \mathbb{R}^{m+nK} with rational coordinates and let \mathbb{Q}_{++} denote the set of strictly positive rational numbers. For a point $(x, r) \in \mathbb{Q}^{m+nK} \times \mathbb{Q}_{++}$, let $K_r(x)$ denote the closed ball in \mathbb{R}^{m+nK} with center x and radius r . For any $y = (f; U_{C^1}) \in \mathcal{N}_F \times \mathcal{N}_{C^1}$, there exists $(x, r) \in \mathbb{Q}^{m+nK} \times \mathbb{Q}_{++}$ with $y \in K_r(x) \subset \mathcal{N}_F \times \mathcal{N}_{C^1}$. Hence the closed set \mathcal{Y} in $\mathcal{N}_F \times \mathcal{N}_{C^1}$ can be covered by a countable family of compact subsets $T_k, k = 1, 2, \dots$, of $\mathcal{N}_F \times \mathcal{N}_{C^1}$, each of the form $T_k = K_r(x) \cap \mathcal{Y}$. For each k , $\hat{T}_k \equiv \Phi(K_k)$ is a compact subset of \mathcal{Z} and $\mathcal{Z} = \Phi(\mathcal{Y}) = \Phi(\bigcup_k T_k) = \bigcup_k \Phi(T_k) = \bigcup_k \hat{T}_k$. Hence \mathcal{Z} is a countable union of compact sets and therefore measurable.

Since each of the sections $Z(U_{C^1})$ has Lebesgue measure zero and \mathcal{Z} is measurable, appli-

cation of Fubini's theorem yields that the graph \mathcal{Z} has Lebesgue measure zero as asserted.

To demonstrate the second assertion, recall the definitions and properties of the open sets $\mathcal{N}_F, \mathcal{N}_{C^2}, \mathcal{N}_{C^1}$, the mapping $\Psi : \mathcal{N}_F \times \mathcal{N}_{C^1} \rightarrow \mathcal{N}_{C^2}$, and the correspondence $Z(\cdot)$. Let $U = (U_{C^1}, U_{C^2}) \in \mathcal{N}_{C^1} \times \mathcal{N}_{C^2} \subseteq \mathbb{R}^{nm}$. Consider further $f \in \mathcal{N}_F$ satisfying $f \in \mathcal{N}_F(U)$ where $\mathcal{N}_F(U)$ is defined in the statement of the Lemma. Suppose also that $U \notin \mathcal{Z}$, i.e. $U_{C^2} \notin Z(U_{C^1})$. Then we claim that f is an isolated point of $\mathcal{N}_F(U)$.

Observe that by Lemma 3, $f \in \mathcal{N}_F(U)$ implies $U_{C^2} = \Psi(f, U_{C^1})$. Because of $U_{C^2} \notin Z(U_{C^1})$, (f, U_{C^1}) is an f -regular point of Ψ , i.e. $\partial\Psi/\partial f$ has full rank at (f, U_{C^1}) . Hence the derivative $D\Psi(\cdot, U_{C^1})$ has full rank at f . Therefore, by the Inverse Function Theorem, there exist open neighborhoods \mathcal{N}_F^* of f in \mathcal{N}_F and $\mathcal{N}_{C^2}^*$ of U_{C^2} in \mathcal{N}_{C^2} such that $\Psi^*(\cdot, U_{C^1})$, the restriction of $\Psi(\cdot, U_{C^1})$ to \mathcal{N}_F^* , is a diffeomorphism from \mathcal{N}_F^* onto $\mathcal{N}_{C^2}^*$.

We conclude that f is the only point in $\mathcal{N}_F^* \cap \mathcal{N}_F(U)$. For suppose $g \in \mathcal{N}_F^* \cap \mathcal{N}_F(U)$. Then $g, f \in \mathcal{N}_F^*$ and $M(g, U) = M(f, U)$ where $U_{C^2} \in \mathcal{N}_{C^2}^*$. Therefore $U_{C^2} = \Psi^*(g, U_{C^1}) = \Psi^*(f, U_{C^1})$. Since $\Psi^*(\cdot, U_{C^1})$ is a diffeomorphism from \mathcal{N}_F^* onto $\mathcal{N}_{C^2}^*$, $g = f$ as asserted. \square

References

- [1] Abreu, D. and A. Rubinstein (1988), "The Structure of Nash Equilibrium in Repeated Games with Finite Automata," *Econometrica*, 56, 1259-1281.
- [2] Anderson, R.M. and W.R. Zame (1997), "Genericity with Infinitely Many Parameters," mimeo, University of California, Berkeley.
- [3] Aumann, R.J. (1981), "Survey of Repeated Games," in Aumann, R.J., et mult. al.: *Essays in Game Theory and Mathematical Economics in Honor of Oskar Morgenstern*, Mannheim: Bibliographisches Institut, 11-42.
- [4] Balasko, Y. (1988), *Foundations of the Theory of General Equilibrium*. Academic Press, Inc. : Boston et al.
- [5] Benabou, R. (1989), "Optimal Price Dynamics and Speculation with a Storable Good," *Econometrica*, 57, 41-80.
- [6] Bergemann, D. and J. Valimaki (1996) "Learning and Strategic Pricing," *Econometrica*, 64, 1125-1149.
- [7] Blume, L.E. and W.R. Zame (1994), "The Algebraic Geometry of Perfect and Sequential Equilibrium," *Econometrica*, 62, 783-794.

- [8] Cho, I.-K. and H. Li (1999), "How Complex are Networks Playing Repeated Games?", *Economic Theory*, 13, 93-123.
- [9] Debreu, G. (1970) "Economies with a Finite Number of Equilibria," *Econometrica*, 38, 387-392.
- [10] Dutta, P. (1995), "A Folk Theorem for Stochastic Games," *Journal of Economic Theory*, 66, 1-32.
- [11] Dutta, P. (1996), "Efficient Markov Perfect Equilibria," mimeo, Columbia University, May.
- [12] Dutta, P. and R. Sundaram (1992), "Markovian Equilibrium in a Class of Stochastic Games: Existence Theorems for Discounted and Undiscounted Models," *Economic Theory*, 2, 197-214.
- [13] Friedman, J.W. (1986), *Game Theory with Applications to Economics*. Oxford University Press: New York and Oxford.
- [14] Fudenberg, D. and E. Maskin (1986), "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information," *Econometrica*, 54, 533-56.
- [15] Fudenberg, D. and E. Maskin (1995), "Learning and Markov Behavior," mimeo, May.
- [16] Gilles, R.P., Haller, H.H., and P.H.M. Ruys (1998), "Semi-Core Equivalence," *Economic Theory*, 11, 101-127.
- [17] Govindan, S., and A. McLennan (1998a), "On the Generic Finiteness of Equilibrium Outcome Distributions in Game Forms," forthcoming in *Econometrica*.
- [18] Govindan, S., and A. McLennan (1998b), "Generic Finiteness of Equilibrium Outcome Distributions for Two person Game Forms with Zero Sum and Common Interest Utilities," mimeo.
- [19] Grandmont, J.-M., Kirman, A.P., and W. Neufeind (1974), "A New Approach to the Uniqueness of Equilibrium", *Review of Economic Studies*, 41, April, 289-291.
- [20] Haller, H. and R. Lagunoff (1997), "Markov Perfect Equilibria in Repeated Asynchronous Choice Games," VPI&SU WP E-97-05, and Economics Working Paper Archive #ewp-game/9707006, <http://econwpa.wustl.edu/eprints/game/papers/9707/9707006.abs> .
- [21] Harsanyi, J.C. (1971), "The Oddness of Equilibrium Points: A New Proof," *International Journal of Game Theory*, 2, 235-250.
- [22] Hildenbrand, W. (1974), *Core and Equilibria of a Large Economy*. Princeton University Press: Princeton.

- [23] Kalai, E., and W. Stanford (1988), "Finite Rationality and Interpersonal Complexity in Repeated Games," *Econometrica*, 56, 397-410.
- [24] Kalai, E., Samet, S., and W. Stanford (1988), "A Note on Reactive Equilibria in the Repeated Prisoner's Dilemma," *International Journal of Game Theory*, 17, 176-186.
- [25] Keiding, H. (1997), "On the Maximal Number of Nash Equilibria in an $n \times n$ Bimatrix Game," *Games and Economic Behavior*, 21, 148-160.
- [26] Kreps, D. and R. Wilson (1982), "Sequential Equilibria," *Econometrica*, 50, 863-894.
- [27] Levhari, D. and L. Mirman (1980), "The Great Fish War: An Example Using a Dynamic Cournot Nash Solution," *Bell Journal of Economics*, 11, 322-334.
- [28] Lagumoff, R. and A. Matsui (1997), "Asynchronous Choice in Repeated Coordination Games," *Econometrica*, 65, 1467-1478.
- [29] Mas-Colell, A. (1985), *The Theory of General Economic Equilibrium: A Differentiable Approach*. Cambridge University Press: Cambridge et al.
- [30] McKelvey, R.D. and A. McLennan (1997), "The Maximal Number of Regular Totally Mixed Nash Equilibria," *Journal of Economic Theory*, 72, 411-425.
- [31] McLennan, A. (1997), "The Maximal Generic Number of Pure Nash Equilibria," *Journal of Economic Theory*, 72, 408-410.
- [32] McLennan, A., and I.-U. Park (1999), "Generic 4×4 Two Person Games Have at Most 15 Nash Equilibria," *Games and Economic Behavior*, 26, 111-130.
- [33] Maskin, E., and J. Tirole (1987), "A Theory of Dynamic Oligopoly, Part III: Cournot Competition", *European Economic Review*, 31, 947-968.
- [34] Maskin, E. and J. Tirole (1988a), "A Theory of Dynamic Oligopoly, I: Overview and Quantity Competition with Large Fixed Costs," *Econometrica*, 56, 549-570.
- [35] Maskin, E. and J. Tirole (1988b), "A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand Curves and Fixed Costs," *Econometrica*, 56, 571- 600.
- [36] Maskin, E. and J. Tirole (1997), "Markov Perfect Equilibrium, I: Observable Actions," mimeo, June.
- [37] Milnor, J.W. (1965), *Topology From the Differentiable Point of View*. The University Press of Virginia: Charlottesville.
- [38] Park, I.-U. (1997), "Generic Finiteness of Equilibrium Outcome Distributions for Sender-Receiver Cheap-Talk Games," *Journal of Economic Theory*, 76, 431-448.
- [39] Quint, T. and M. Shubik (1997), "A Theorem on the Number of Nash Equilibria in a Bimatrix Game," *International Journal of Game Theory*, 26, 353-359.

- [40] Rosenmüller, J. (1971), “On a Generalization of the Lemke-Howson Algorithm to Non-cooperative N -Person Games,” *SIAM Journal of Applied Mathematics*, 21, 73-79.
- [41] Shapley, L. (1953), “Stochastic Games,” *Proceedings of the National Academy of Sciences*, 39, 1095-1100.
- [42] Wilson, R. (1971), “Computing Equilibria in N -Person Games,” *SIAM Journal of Applied Mathematics*, 21, 80-87.