

Markov Perfect Equilibria in Repeated Asynchronous Choice Games

Hans Haller* and Roger Lagunoff†

December 4, 2005‡

Abstract

This paper examines the issue of multiplicity of Markov Perfect equilibria in alternating move repeated games. Such games are canonical models of environments with repeated, asynchronous choices due to inertia or replacement. Our main result is that the number of Markov Perfect equilibria is generically finite with respect to stage game payoffs. This holds despite the fact that the stochastic game representation of the alternating move repeated game is “non-generic” in the larger space of state dependent payoffs.

Journal of Economic Literature Classification Numbers: C72, C73.

Keywords: repeated games, asynchronous choice, alternating moves, stochastic games, Markov Perfect equilibria, genericity.

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

†Department of Economics, Georgetown University, Washington DC 20057 USA. 202-687-1510, lagunofr@georgetown.edu.

‡We would like to thank, without implicating, Prajit Dutta, Drew Fudenberg, Daniel Kahn, Akihiko Matsui, Andrew McLennan, Lones Smith, Bernard Cornet, and two anonymous referees for helpful comments and suggestions.

1 Introduction

There are many repeated game settings in which players move sequentially or *asynchronously*. For example, capacity decisions in an oligopolistic industry which involve some inertia or costly upgrading are difficult to synchronize. Maskin and Tirole (1988a) consider a model in which each firm sets capacity given the (temporarily) fixed capacity of rivals.¹ Birth/death processes in evolutionary models provide another example. Replacements occur at independent, random times, hence are generically asynchronous.

Since most asynchronously repeated games, including the alternating move game, are stochastic games, Folk Theorems of Dutta (1995) and later Yoon (2001) show that if a full dimensionality condition holds, then every stage game payoff above each player's effective minmax can be implemented by Subgame Perfect equilibrium when the players are sufficiently patient. Hence, asynchronous timing does not, by itself, reduce the multiplicity of equilibria.²

The present paper takes a different route by examining multiplicity of equilibria in alternating move games when the strategy profiles are Markov Perfect equilibria (MPE). MPE are Subgame Perfect equilibria in Markov strategies. The use of Markov strategies rules out the possibility to extend memory beyond what is encoded in the states of the game. MPE is a natural solution concept when players find it difficult or costly to coordinate their actions on the past history of play. For example, Bhaskar and Vega Redondo (2002) show that any Subgame Perfect equilibrium of the alternating move game in which players' memory is bounded and their payoffs reflect the costs of strategic complexity must coincide with a MPE.

Generally, Markov Perfect equilibria in games with alternating moves are different than in games with simultaneous moves. In the latter case, MPE are trivial. With simultaneous moves, there are no non-trivial state variables, and so Markov Perfect equilibria are merely repetitions of Nash equilibria of the stage game.³ In the case of alternating moves, however, non-trivial payoffs (different from stage game equilibrium payoffs) are sustainable. In Section 3, we show that MPE can support mutual cooperation in some Prisoner's Dilemma games. Therefore, while the restriction to Markov Perfect equilibria pares down the equilibrium set in both the asynchronous and the simultaneous move repeated game, the state variables in the asynchronous model admit enough flexibility to support desirable outcomes which could not arise in MPE of the simultaneous move game.

¹See also Maskin and Tirole (1987, 1988b).

²Lagunoff and Matsui (1997) prove a uniqueness result in the special case in which the stage game is a pure coordination game (a game which violates the full dimensionality restriction).

³In that case, well known results of Rosenmüller (1971), Wilson (1971) and Harsanyi (1973) establish that the number of Nash equilibria in any normal form (stage game) is generically finite with respect to payoffs.

Hence, how “much” multiplicity is there with regard to Markov Perfect equilibria of asynchronous choice games? A recent result of Haller and Lagunoff (2000) (HL from here on) shows that Markov Perfect equilibria are generically finite in the state contingent payoffs of the stochastic game.⁴ The HL theorem would appear to apply to asynchronous choice games such as alternating moves. Yet, *it does not*.

To see why, notice that HL’s result applies to all stochastic games with n players, k states, and $\ell(j)$ action profiles in state $j = 1, \dots, k$. Letting $q = \sum_{j=1}^k \ell(j)$, the HL result shows that for all games except possibly a set of stochastic game payoff vectors with Lebesgue measure 0 in \mathbb{R}^{nq} , the set of MPE is finite.

The problem in terms of the present model, is that all alternating move games are contained in a Lebesgue measure 0 subset of \mathbb{R}^{nq} ! Consider, for example, the alternating move game derived from the stage game in Figure 1.

		2	
		L	R
	T	5, 3	20, 2
1	M	10, 1	0, 4
	B	11, 6	9, 18

Figure 1

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. In an odd period, if, for example, player 1 takes T and if player 2 took L in the previous period, then the realized payoff of this period to player $i = 2$ is $u_2(T, L) = 3$. Payoffs in even periods are similarly defined. The state space is the set $\{T, M, B, L, R\}$, corresponding to possible actions taken in the previous period. If the current state is R , for example, then this means that player 2 moved in the previous period and took action R . A Markov strategy for this game is a mapping for each player from states to mixed strategies.

⁴Subsequently, Herings and Peeters (2004) applied homotopy theory to select and compute MPE. As a by-product, they obtained that generically, the number of MPE is odd.




Figure 2

Figure 2 demonstrates the standard stochastic game representation of alternating moves. The arrows represent potential transitions. The payoffs in this game represent a point in \mathbb{R}^{24} . The “full dimensional” Genericity Theorem of HL asserts that on a set of full Lebesgue measure in \mathbb{R}^{24} , there are finitely many MPE of stochastic games of the form in Figure 2. A key to that result, it turned out, is the ability to rewrite stochastic game payoffs as a linear function of the 24×1 state-dependent payoff vector of the stochastic game. This facilitates the use of standard perturbation techniques to prove the genericity result.

Unfortunately, that result does not apply to asynchronous move games. 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.⁵

This paper therefore proves a lower dimensional genericity theorem. We show that the set of Markov Perfect equilibria for the alternating move is generically finite in the set of *stage game payoffs*. Applying this paper’s result to the example would show indeed that there are finitely many MPE on a generic set in \mathbb{R}^{12} , rather than in \mathbb{R}^{24} .

The key to the present result is to encode payoffs in such a way that each Markov strategy profile induces a Markov transition matrix that has lower dimension than in the standard

⁵There is an alternative “standard representation” of a stochastic game in which action spaces are invariant across states by creating redundant moves. Of course, this makes the genericity problem even worse. In the present case, Player 1 would have 3 feasible actions, and Player 2 two feasible actions in each state. The state dependent payoff vectors would then reside in \mathbb{R}^{60} since there are 12 possible stage game payoffs and 5 possible states.

stochastic game formulation. The encoding entails that states be identified with stage game strategies, and, consequently, asynchronous move game payoffs may be expressed as linear functions of the stage game payoff vector. With this encoding, a modification of standard generic finiteness techniques of Debreu (1970) can be applied.

Apparently, a crucial feature of the alternating move game is that stage game actions play a dual role since they also define the current state. For this reason, this encoding does not extend to other stochastic games. Nevertheless, we hope the current result might yield some insight toward a general “lower dimensional” genericity theorem with applications to other nongeneric yet useful stochastic games. Exactly how to proceed is an open question.⁶

An extension of the current result to asynchronous games with more than two players is straightforward. But we refrain from doing so, because there are numerous ways to write down asynchronous models with $n > 2$ players and the general the n -player asynchronous game is notationally intensive.

The balance of this paper is organized as follows. Section 2 defines the canonical alternating move game, defines MPE, and states the main result. Section 3 gives an example of a 2×2 in which we show that cooperation can be sustained in MPE if the game is a certain Prisoner’s Dilemma game. We also show that completely mixed MPE cannot hold generically, providing some useful intuition for the theorem. Section 4 elaborates on an alternative specification of payoffs that facilitates the proof. Section 5 gives the proof. Some of the details of the proof (i.e., proofs of the lemmata) are contained in the Appendix, Section 6.

2 The Canonical Model

2.1 Preliminaries

Consider the alternating move game associated with the stage game given as $G = \langle S_1, S_2, u_1, u_2 \rangle$ where S_i , $i = 1, 2$, is a finite set of actions of player i , and, letting $S = S_1 \times S_2$, $u_i : S \rightarrow \mathbb{R}$ is the utility function of player i . Alternatively, let U_i denote i ’s $|S| \times 1$ utility vector so that $U_{is} = u_i(s)$. Then $U = (U_1, U_2)$ is an element of \mathbb{R}^{2S} .

⁶This explains, in part, why our attempts at a meta-theorem encompassing both this and the HL (2000) result have failed so far. more than 2 players. We think an extension is straightforward, but did not prove it for 2 reasons. First, there are numerous ways to write down asynchronous model with $n > 2$ and the general the n -Player asynchronous game is notationally intensive. Second, it would lengthen the paper without adding much insight, since the logic applies directly. For other types of games and equilibrium concepts, some lower dimensional genericity results exist. Examples are Govindan and McLennan (1999) for extensive game forms, and Park (1997) for signalling games.

After the first decision node, which occurs for all players at time zero, the two players alternately have chances to revise their actions. 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. In an odd period, if player 1 takes s_1 and if player 2 took s_2 in the previous period, then the realized payoff of this period to player $i = 1, 2$ is $u_i(s_1, s_2)$. Payoffs in even periods are similarly defined. Players have common discount factor δ with $0 \leq \delta < 1$.

Let $s^t = (s_1^t, s_2^t)$ denote the profile of actions in period t . As in ordinary repeated games, individuals seek to maximize the discounted sum $\sum_{t=0}^{\infty} (1 - \delta)\delta^t u_i(s^t)$. Let H_1 denote the set of all histories ending in odd numbered periods and H_2 denote the set of all histories ending in even numbered periods. Let $H = H_1 \cup H_2$. A standard notation denotes the history ending in period t by h^t . A *strategy* for player $i = 1, 2$ is a function $f_i : H_j \rightarrow \Delta(S_i)$ where $\Delta(S_i)$ denote the set of mixed strategies on S_i . We write $f_i(s_i|h)$ to denote the probability weight assigned s_i when the current history is h . A *Subgame Perfect Equilibrium* $f^* = (f_1^*, f_2^*)$ is a strategy profile in which for each $i = 1, 2$, f_i^* is a best response to f_j^* , $j \neq i$, after every history $h \in H$.

We restrict our attention to the special class of Perfect Equilibria known as Markovian or *Markov Perfect Equilibria* (MPE). A Markov Perfect Equilibrium is a Perfect Equilibrium in *Markov strategies*, that is, strategies that depend only on payoff relevant information. In the alternating move game the payoff relevant state at time t is the action s_j^{t-1} of the previous mover j at time $t - 1$. It will later prove useful to distinguish states from actions, although they are in one-to-one correspondence. Formally, a Markovian strategy can be expressed (abusing notation somewhat) as a strategy $f_i : S_j \rightarrow \Delta S_i$. We use the notation $f_j(s_j|s_i)$ to denote the probability that j assigns to s_j given that the current state (the action taken by the previous mover) is s_i .

Rewriting the expression for payoffs now as a function of only the state rather than the entire history gives for $i = 1$ the recursive expression,

$$V_1(f|s_i) = \sum_{s_j} f_j(s_j|s_i) [(1 - \delta)u_1(s_i, s_j) + \delta V_1(f|s_j)] \quad (1)$$

A similar expression for $i = 2$ is obtained by switching subscripts.

A *Markov Perfect equilibrium* $f^* = (f_1^*, f_2^*)$ is a strategy profile in which for each $i = 1, 2$, after every state $s_j \in S_j$,

$$V_i(f^*|s_j) \geq V_i(f_i, f_j^*|s_j)$$

for any possibly non-Markov strategy f_i . Clearly, a *Markov Perfect Equilibrium* is a Subgame Perfect equilibrium in which the strategies are Markovian. Existence of Markov Perfect Equilibria is a standard result. See, for example, Friedman (1986) and Fudenberg and Tirole (1991).

2.2 Main Result

The main result of this paper is:

Theorem *The set of MPE is a finite set on a full measure subset of $U \in \mathbb{R}^{2|S|}$.*

Notice, first, that the conclusion does not depend on the discount factor. 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. Second, notice that, unlike in normal form games, strategies do not enter linearly in payoffs. Since players' payoffs are expected discounted sums of stage game payoffs, players' payoffs are shown to be rational functions of Markovian strategies. For this reason, a direct application of Harsanyi's (1973) proof is not possible. We elaborate on the logic used in the proof in Section 4.

3 An Example

We illustrate the logic of the result as well as some key differences between MPE of asynchronous and simultaneous play. Consider the stage game below.

		2	
		<i>C</i>	<i>D</i>
1	<i>C</i>	a_1, a_2	b_1, d_2
	<i>D</i>	d_1, b_2	c_1, c_2

Figure 3

We examine some properties of MPE of the alternating move version of the stage game in Figure 3. Fix a strategy profile f and let $\beta_i = f_i(D|D)$ denote the probability that each player plays D whenever the other did. Let $\alpha_i = f_i(C|C)$ denote the probability that C is played by i after j chose C .

Let V_{iC} denote the value of player i 's randomization when it is his turn to move and player j chose C in the previous period. Let W_{iC} denotes the value to i when it is the other player's turn to move and i chose C previously. Define V_{iD} and W_{iD} similarly. Using the expression for payoffs in (1), observe that

$$V_{iC} = \alpha_i ((1 - \delta)a_i + \delta W_{iC}) + (1 - \alpha_i) ((1 - \delta)d_i + \delta W_{iD}) \quad (2)$$

and

$$V_{iD} = \beta_i((1 - \delta)c_i + \delta W_{iD}) + (1 - \beta_i)((1 - \delta)b_i + \delta W_{iC}) \quad (3)$$

Example 1. Prisoner's Dilemma. Suppose the stage game is a Prisoner's Dilemma game, i.e., $d_i > a_i > c_i > b_i$ for each i , and suppose $a_i + c_i > b_i + d_i$ for each i . Then, unlike in simultaneous move games, we show that there is a MPE that sustains long run average payoffs arbitrarily close to a_i if δ close to one.

To see this, consider the following strategy. Let $\alpha_i = 1$ for both i . Each player chooses C for sure when the other chose C previously. For each i and each $j \neq i$, let

$$\beta_j = \frac{(c_i + d_i) - (a_i + b_i)}{(a_i + d_i) - (c_i + b_i)}$$

It is easy to verify that $0 < \beta_j < 1$ for each j . To verify that this Markovian strategy is an MPE, we verify that it satisfies the three conditions which characterize MPE converging globally to (C, C) :

$$a_i > (1 - \delta)d_i + \delta W_{iD} \quad (4)$$

$$V_{iD} = (1 - \delta)c_i + \delta W_{iD} = (1 - \delta)b_i + \delta a \quad (5)$$

$$W_{iD} = \beta_j[(1 - \delta)c_i + \delta V_{iD}] + (1 - \beta_j)[(1 - \delta)d_i + \delta a] \quad (6)$$

The first inequality defines the incentive constraint to remain in (C, C) once it is reached. The second is the equality constraint for using a mixed strategy β_i . The last is the payoff after choosing D when the other player will use his mixed strategy β_j . If $a_i + c_i > b_i + d_i$ then all three conditions are satisfied for the MPE which we constructed. Under this strategy, all payoffs above converge to a_i as $\delta \rightarrow 1$.

Example 2. No Generically Mixed MPE. Finally, we show that strictly mixed MPE are generically impossible. That is, except on a set of payoffs (a_i, b_i, c_i, d_i) with measure zero in \mathbb{R}^4 , there are no MPE in which any player uses a strictly mixed strategy, i.e., there are no MPE in which $0 < \alpha_i < 1$ and $0 < \beta_i < 1$ for any i .

If $0 < \alpha_i < 1$, i.e., i strictly randomizes after j chose C , then a standard property of equilibrium is that:

$$V_{iC} = (1 - \delta)a_i + \delta W_{iC} = (1 - \delta)d_i + \delta W_{iD}, \quad (7)$$

since i must be indifferent between C and D in state C . Similarly, if $0 < \beta_i < 1$, i.e., player i strictly randomizes after j chose D then

$$V_{iD} = (1 - \delta)c_i + \delta W_{iD} = (1 - \delta)b_i + \delta W_{iC}. \quad (8)$$

However, equations (7) and (8) are mutually consistent only if $d_i + b_i = a_i + c_i$. The 4-tuples of payoffs for which $d_i + b_i = a_i + c_i$ spans a 3-dimensional hyperplane in \mathbb{R}^4 .

4 Payoff Representation

In this example, dynamic payoffs V are linear functions of stage payoffs. It is not hard to show that this is generally the case. Generally, an alternative to the expression (1) for payoffs will prove more convenient. Observe that any Markovian profile f may be expressed as a Markov chain on the space of profiles S . Let $\mathcal{P}_f^i, i = 1, 2$ denote an $S \times S$ transition matrix in which $f_i(s'_i | s_j)$ denotes the element corresponding to row s and column s' . The rows are then indexed by the current states, while the columns are indexed by subsequent states. For each profile \hat{s} define the set of *player i 's adjacent states of \hat{s}* by:

$$\Theta_i(\hat{s}) \equiv \{s \in S \mid \hat{s}_j = s_j\}$$

In words, the adjacent states are those that can reach s by player i 's unilateral decision. Clearly $s \in \Theta_i(s')$ iff $s' \in \Theta_i(s)$. It is clear that if \mathcal{P}_f^i is any transition matrix derived from a strategy f_i , then for any two adjacent profiles s, s' we must have for any column of the matrix that the entries corresponding to the two rows s and s' are equal.

As an example, suppose that in the stage game above in Figure 3, both players choose a deterministic Markovian strategy in which each player chooses C given that the other chose C , player 2 chooses D after 1 chose D , and player 1 chooses C after 2 chose D . In this case, the Markov chain may be expressed as the arrows of a directed graph as seen in Figure 4a below. The corresponding transition matrix for each player is given by Figure 4b.




Figure 4a

$$\mathcal{P}_f^1 = \begin{array}{c} \begin{array}{cccc} a & d & c & b \end{array} \\ \begin{bmatrix} 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \end{bmatrix} \end{array}, \quad \mathcal{P}_f^2 = \begin{array}{c} \begin{array}{cccc} a & d & c & b \end{array} \\ \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 \end{bmatrix} \end{array}$$

Figure 4b

Let π denote any $1 \times S$ distribution of initial profiles $s = s^0$. Let $\pi(s)$ denote the value of the s th component of π , and let π_s denote the particular (Dirac or unit) vector which

assigns 1 to component $\pi_s(s)$ and 0 elsewhere. Then i 's payoff of any Markovian profile f when it is i 's turn to move given j 's choice of s_j in the previous period is given by:

$$V_i(f|s_j) = \pi_s \cdot \left[\sum_{t=0}^{\infty} (1-\delta)\delta^{2t} [\mathcal{P}_f^i \cdot \mathcal{P}_f^j]^t \cdot \mathcal{P}_f^i \cdot [I + \delta\mathcal{P}_f^j] \right] \cdot U_i \quad (9)$$

where I is the $S \times S$ identity matrix, π_s is the initial distribution placing unit mass on initial state s , and U_i denotes the $S \times 1$ vector of utilities over profiles $s \in S$. Let

$$\begin{aligned} A_f^i &= \left[\sum_{t=0}^{\infty} (1-\delta)\delta^{2t} [\mathcal{P}_f^i \cdot \mathcal{P}_f^j]^t \cdot \mathcal{P}_f^i \cdot [I + \delta\mathcal{P}_f^j] \right] \\ &= (1-\delta)[I - \delta^2(\mathcal{P}_f^i \cdot \mathcal{P}_f^j)]^{-1} \cdot \mathcal{P}_f^i \cdot [I + \delta\mathcal{P}_f^j] \end{aligned} \quad (10)$$

where the inverse $[I - \delta^2(\mathcal{P}_f^i \cdot \mathcal{P}_f^j)]^{-1}$ exists because $[I - \delta^2(\mathcal{P}_f^i \cdot \mathcal{P}_f^j)]$ is a matrix with a strictly dominant diagonal. Equation (9) can now be rewritten as

$$V_i(f|s_j) = \pi_s \cdot A_f \cdot U_i. \quad (11)$$

From (11) it is clear that although the dynamic payoff is not multilinear in strategy profiles, f , it is linear in stage game payoffs U_i . The proof of the main result utilizes this key observation. It follows then that a certain Jacobian matrix derived from the first order conditions of (11) is linear in stage game payoffs. After fixing supports (carriers) and normalizations of Markovian profiles and narrowing the domain of stage game payoff matrices first to small neighborhoods and further to sections of those, we arrive at a local and lower-dimensional genericity result, using the Implicit Function Theorem and Sard's Theorem. Application of Fubini's Theorem helps lift the result to higher dimensions. A countability argument turns the local result into a global one. While the conclusion of Sard's Theorem is about regular values, it can be converted into a statement about locally isolated points which, combined with compactness, yields finiteness. The possible co-existence of MPE with different carriers complicates, but does not invalidate the argument.

5 The Proof

5.1 Notation

We begin by introducing some necessary notation. Let $F = F_1 \times F_2$ denote the set of all Markov profiles. Clearly, F is a compact and convex subset of $\mathbb{R}^{2|S|}$. Let $E(U)$ denote the set of MPE given payoff vector U . Since the MPE are defined by weak inequalities, and since

payoffs, as given by (11), are smooth functions of f (as well as U_i), it follows that $E(U)$ is a closed, and therefore compact subset of F .

We follow Harsanyi (1973) and partition the set F with respect to carriers or supports. The *support* of $f_i \in F_i$ is defined as the set $C_i(f_i) = \{s \in S : f_i(s_i|s_j) \neq 0\}$. A joint Markov strategy $f = (f_1, f_2) \in F$ has support $C(f) = (C_1(f_1), C_2(f_2))$.

Given a fixed support $C = (C_1, C_2)$, let F_C denote the set of $f \in F$ with support $C = (C_1, C_2)$, i.e., $f \in F_C$ if and only if $C = C(f)$. Observe that F_C is the interior of a face the simplex F , hence is an open subset of $\mathbb{R}^{|C_1|} \times \mathbb{R}^{|C_2|}$. Let $E_C(U)$ denote the set of MPE with support C .

Observe that

$$\bigcup_C F_C = F \quad \text{and} \quad \bigcup_C E_C(U) = E(U) \quad (12)$$

For the remainder of the proof, we *fix* a support, C , and then proceed to show that $E_C(U)$ is a finite set on a full measure subset of $U \in \mathbb{R}^{2|S|}$. Since there are only finitely many supports, the conclusion of the Theorem must then follow from (12).

Hence, fix a support $C = (C_1, C_2)$. Next, we fix a normalization for this support. Specifically, for each i and $s_j \in S_j$, we fix some action $d(s_j) \in S_i$ for Player i . Let $D_i = \{(d(s_j), s_j) : s_j \in S_j\}$ and $D = (D_1, D_2)$. Notice that $|D_1| = |S_2|$ and $|D_2| = |S_1|$. Varying $s_j \in S_j$ for $j \neq i$, $i = 1, 2$, we pin down $|S_1| + |S_2|$ normalizing equations satisfying

$$f_i(d(s_j)|s_j) = 1 - \sum_{s_i \neq d(s_j)} f_i(s_i|s_j), \quad d(s_j) \in S_i, \quad s_j \in S_j, \quad j \neq i, \quad i = 1, 2. \quad (13)$$

Let $R_i = C_i \setminus D_i$. Notice that $R_i = \emptyset$ conforms to the special case in which Player i always chooses a pure strategy in every state. Clearly the set of pure strategy MPE must be finite. The case of, say, $R_1 = \emptyset$ and $R_2 \neq \emptyset$ can also be handled analogously to the case of $R_i \neq \emptyset$ for both $i = 1$ and $i = 2$. Hence, we restrict attention to supports in which $R_i \neq \emptyset$ for both i . Notice $|R_1| = |C_1| - |S_2|$ and $|R_2| = |C_2| - |S_1|$.

For each $i = 1, 2$, define the mapping $J_i : F_C \times \mathbb{R}^{|S|} \rightarrow \mathbb{R}^{|R_i|}$ by: for all $s \in R_i$,

$$J_{i s}(f, U_i) = \pi_s \cdot B_f(s_i|s_j) \cdot U_i \quad (14)$$

where $B_f(s_i|s_j)$ is the $S \times S$ matrix defined by

$$B_f(s_i|s_j) = \frac{\partial}{\partial f_i(s_i|s_j)} [A_f^i]$$

A few comments are in order. First, J_i is simply the Jacobian map of the payoff function V_i on the support C . Because $f_i(s_i|s_j)$ is just an interior probability number, the

partial $\frac{\partial}{\partial f_i(s_i|s_j)} [\cdot]$ is well defined. Second, using the normalization equations in (13) above, $f_i(s_i|s_j)$ also appears when $f_i(d(s_j)|s_j)$ appears in the matrix $B_f(s_i|s_j)$. Third, by varying the $f_i(s_i|s_j)$, every term in the matrix is potentially affected.

5.2 The Steps of the Proof

We now prove the Theorem using the following Steps or Lemmata. The Proof of each Lemma is contained in the Appendix.

The first, and perhaps most critical step in the argument is:

Lemma 1 *Let $f \in F_C$. Then for all $U_i \in \mathbb{R}^{|S|}$,*

$$D_{U_i} J_i(f, U_i) \text{ has full rank of } |R_i|.$$

Lemma 1 is a key step since it establishes full rank with respect to the “lower dimensional” set R_i .

Observe that for the mapping J defined by $J(f, U) \equiv (J_1(f, U_1), J_2(f, U_2))$, it clear that $D_U J$ is of the block matrix form

$$D_U J(f, U) = \begin{bmatrix} D_{U_1} J_1(f, U_1) & 0 \\ 0 & D_{U_2} J_2(f, U_2) \end{bmatrix}$$

Figure 5

Consequently, Lemma 1 implies that $D_U J$ has full rank of $|R_1| + |R_2|$.

Lemma 2 *Fix $(f^*, U^*) \in F_C \times \mathbb{R}^{2|S|}$ such that $J(f^*, U^*) = 0$. Then, for each i there exists a partition, (P_i, Q_i) , of C_i (i.e., $P_i \cup Q_i = C_i$ and $P_i \cap Q_i = \emptyset$) with the following properties:*

1. $|P_i| = |R_i|$ and $|Q_i| = |C_i| - |R_i|$ with corresponding notation for subvectors of U^* ,

$$U_P^* \equiv (U_{P_1}^*, U_{P_2}^*) \quad \text{and} \quad U_Q^* \equiv (U_{Q_1}^*, U_{Q_2}^*),$$

2. there are open sets \mathcal{F} , \mathcal{U}_P and \mathcal{U}_Q that contain f^* , U_P^* and U_Q^* , resp., and a locally smooth mapping $\Psi : \mathcal{F} \times \mathcal{U}_Q \rightarrow \mathcal{U}_P \times \mathcal{U}_Q$ such that for all $(f, U_Q) \in \mathcal{F} \times \mathcal{U}_Q$, the vector $(U_P, U_Q) = \Psi(f, U_Q)$ is the unique solution to $D_U J(f, (U_P, U_Q)) = 0$.

As before, the proof is in the Appendix. For the next two lemmata, fix (f^*, U^*) satisfying the properties of Lemma 2, and fix the corresponding neighborhoods \mathcal{F} , \mathcal{U}_P and \mathcal{U}_Q and the smooth mapping Ψ .

Call $(f, U_Q) \in \mathcal{F} \times \mathcal{U}_Q$ an *f-regular point* of Ψ , if $\partial\Psi/\partial f$ has full rank $|R_1|+|R_2|$ at (f, U_Q) . Call (f, U_Q) an *f-critical point* of Ψ otherwise. Call $(U_P, U_Q) \in \mathcal{U}_P \times \mathcal{U}_Q$ an *f-critical value* of Ψ , if $(U_P, U_Q) = \Psi(f, U_Q)$ with (f, U_Q) an *f-critical point* of Ψ .

Lemma 3 *The set of f-critical values of Ψ has Lebesgue measure zero in $\mathcal{U}_P \times \mathcal{U}_Q$.*

Clearly, then, the set of *f-regular values* has full Lebesgue measure in $\mathcal{U}_P \times \mathcal{U}_Q$.

Lemma 4 *Let U be an f-regular value of Ψ . Then, the set*

$$\{f \in \mathcal{F} : J(f, U) = 0\}$$

consists of isolated points.

The foregoing definition and analysis of the mapping J can be extended to a domain $G_C \times \mathbb{R}^{|S|}$ where G_C is open and contains the closure of F_C , because $\delta \in [0, 1)$ and the matrices \mathcal{P}_f^1 and \mathcal{P}_f^2 are stochastic. Let us fix such a G_C .

Without loss of generality, we may assume that a set as above \mathcal{F} is an open box with rational-valued end points. Thus, by varying over all pairs (f^*, U^*) with $J(f^*, U^*) = 0$ and suitable partitions (P, Q) , we obtain a countable collection of open sets $\{\mathcal{F}^n\}_{n=1}^\infty$ and a corresponding sequence of locally smooth functions $\{\Psi^n\}$ and open sets $\{\mathcal{U}_Q^n\}$ and $\{\mathcal{U}_P^n\}$, with $\Psi^n : \mathcal{F}^n \times \mathcal{U}_Q^n \rightarrow \mathcal{U}_Q^n \times \mathcal{U}_P^n$ such that:

- (a) The set

$$\mathcal{W}_C \equiv \mathbb{R}^{2|S|} \setminus \left(\bigcup_n \{U \in \mathcal{U}_Q^n \times \mathcal{U}_P^n : U \text{ is an } f\text{-critical value of } \Psi^n\} \right)$$

has full Lebesgue measure in $\mathbb{R}^{2|S|}$; and

(b) for each $U \in \mathcal{W}_C$ (with \mathcal{W}_C defined as in (a)), the set

$$\mathcal{X}(U) \equiv \{f \in G_C : J(f, U) = 0\}$$

consists of isolated points, and each $f \in \mathcal{X}(U)$ is contained in some \mathcal{F}^n .

Because all points in $\mathcal{X}(U)$ are locally isolated, without loss of generality we take the open sets \mathcal{F}^n sufficiently small so that

$$|\mathcal{F}^n \cap \mathcal{X}(U)| = 1.$$

In other words, each such $f \in \mathcal{X}(U)$ is contained in a distinct \mathcal{F}^n .

Lemma 5 *For each $U \in \mathcal{W}_C$, the set $\mathcal{X}(U) \cap F_C$ is finite.*

Lastly, we establish:

Lemma 6 *For each $U \in \mathcal{W}_C$, $E_C(U) \subseteq \mathcal{X}(U) \cap F_C$.*

By combining Lemma 5 with Lemma 6, it follows that for each $U \in \mathcal{W}_C$, the set $E_C(U)$ is a finite set. Finally, recalling that $\bigcup_C E_C(U) = E(U)$ with the union over finitely many supports, it follows that $E(U)$ is a finite set whenever $U \in \mathcal{W} = \bigcap_C \mathcal{W}_C$. Since each of the finitely many sets \mathcal{W}_C has full Lebesgue measure, so has \mathcal{W} , which concludes the proof. ■■

6 Appendix

Proof of Lemma 1

Observe that each entry in $D_{U_i} J_i(f, U_i)$ assumes the form

$$\frac{\partial J_{i s}(f, U_i)}{\partial U_{i s'}} = \pi_s \cdot B_f(s_i | s_j) \cdot 1_{s'} \quad (15)$$

where $s = (s_i, s_j) \in R_i$ and $1_{s'}$ is the $|S| \times 1$ Dirac vector with a 1 in the s' position and zeroes elsewhere.

Observe further that

$$\begin{aligned}
B_f(s_i|s_j) &= \frac{\partial}{\partial f_i(s_i|s_j)} [A_f^i] \\
&= \frac{\partial}{\partial f_i(s_i|s_j)} [(1-\delta)[I - \delta^2 \mathcal{P}_f^i \cdot \mathcal{P}_f^j]^{-1} \cdot \mathcal{P}_f^i \cdot [I + \delta \mathcal{P}_f^j]] \\
&= (1-\delta) \left([I - \delta^2 \mathcal{P}_f^i \cdot \mathcal{P}_f^j]^{-1} \cdot \left[\frac{\partial \mathcal{P}_f^i}{\partial f_i(s_i|s_j)} \right] \cdot [I + \delta \mathcal{P}_f^j] \right. \\
&\quad \left. - [I - \delta^2 \mathcal{P}_f^i \cdot \mathcal{P}_f^j]^{-1} \right. \\
&\quad \left. \cdot \left[-\delta^2 \frac{\partial \mathcal{P}_f^i}{\partial f_i(s_i|s_j)} \cdot \mathcal{P}_f^j \right] \cdot [I - \delta^2 \mathcal{P}_f^i \cdot \mathcal{P}_f^j]^{-1} \cdot \mathcal{P}_f^i \cdot [I + \delta \mathcal{P}_f^j] \right) \\
&= (1-\delta) [I - \delta^2 \mathcal{P}_f^i \cdot \mathcal{P}_f^j]^{-1} \cdot \left[\frac{\partial \mathcal{P}_f^i}{\partial f_i(s_i|s_j)} \right] \\
&\quad \cdot [I + \delta^2 \mathcal{P}_f^j \cdot [I - \delta^2 \mathcal{P}_f^i \cdot \mathcal{P}_f^j]^{-1} \cdot \mathcal{P}_f^i] \cdot [I + \delta \mathcal{P}_f^j].
\end{aligned} \tag{16}$$

In the definition of $B_f(s_i|s_j)$, $\frac{\partial \mathcal{P}_f^i}{\partial f_i(s_i|s_j)}$ denotes the entry by entry derivative of \mathcal{P}_f^i with respect to $f_i(s_i|s_j)$. The term $\frac{\partial \mathcal{P}_f^i}{\partial f_i(s_i|s_j)}$ is a $|S| \times |S|$ matrix that assigns value 1 to the entries in \mathcal{P}_f^i in column $s = (s_i, s_j)$ in which the row profiles s' are adjacent to profile s , i.e., those rows s' in which $s' \in \Theta_i(s)$. It also assigns value -1 to entries of column $(d(s_j), s_j)$ which are adjacent to profile s (which are also, by construction, adjacent to profile $(d(s_j), s_j)$). Zeroes occur everywhere else in the matrix. Hence, this matrix has the form:

$$\frac{\partial \mathcal{P}_f^i}{\partial f_i(s_i|s_j)} = \begin{bmatrix} 0 & 0 & \cdots & \cdots & \cdots & 0 & 0 \\ 0 & 0 & \cdots & \cdots & \cdots & 0 & 0 \\ 0 & 1 & 0 & \cdots & 0 & -1 & 0 \\ 0 & 1 & 0 & \cdots & 0 & -1 & 0 \\ \vdots & \vdots & \vdots & & \vdots & \vdots & \vdots \\ 0 & 1 & 0 & \cdots & 0 & -1 & 0 \\ 0 & 0 & \cdots & \cdots & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & & \vdots & \vdots & \vdots \\ 0 & 0 & \cdots & \cdots & \cdots & 0 & 0 \end{bmatrix}$$

Figure 6

Let $\Lambda_s(i)$ denote the $1 \times |S|$ row vector $\pi_s \cdot [I - \delta^2 \mathcal{P}_f^i \cdot \mathcal{P}_f^j]^{-1} \cdot [\frac{\partial \mathcal{P}_f^i}{\partial f_i(s_i|s_j)}]$. Now let $a_{s's''}$ denote an entry in matrix $[I - \delta^2 \mathcal{P}_f^i \cdot \mathcal{P}_f^j]^{-1} = \sum_t \delta^{2t} (\mathcal{P}_f^i \cdot \mathcal{P}_f^j)^t$. Given the form taken by matrix $\frac{\partial \mathcal{P}_f^i}{\partial f_i(s_i|s_j)}$ in Figure 6 one can verify that

$$\Lambda_s(i) = \left(0, \dots, 0, \sum_{s' \in \Theta_i(s)} a_{s's}, 0, \dots, 0, - \sum_{s' \in \Theta_i(s)} a_{s's}, 0, \dots, 0 \right) \quad (17)$$

with potential nonzero entries in the s th and the $(d(s_j), s_j)$ th position.

We assert that $\sum_{s' \in \Theta_i(s)} a_{s's} \neq 0$. To see why, observe that the entries $a_{s's''}$ are nonnegative since they are discounted sums and products of probability numbers. Moreover, $a_{ss} > 0$ since the first term in the sum $\sum_{t=0}^{\infty} \delta^{2t} (\mathcal{P}_f^i \cdot \mathcal{P}_f^j)^t$ is the identity matrix I . Finally, recall that $\Theta_i(s)$ is the set of profiles s' reached by player i 's unilateral departure from s in the stage game, and so $s \in \Theta_i(s)$ trivially. Hence, $\sum_{s' \in \Theta_i(s)} a_{s's} \neq 0$.

Therefore, any linear combination $\sum_s \lambda_s \Lambda_s(i)$ for a nonzero weight vector $\lambda = (\lambda_s)$ cannot be the zero vector. This means that the $|R_i| \times |S|$ matrix $\Lambda(i) \equiv [\Lambda_s(i)]_{s \in R_i}^T$ has full row rank of $|R_i|$.

Now since

$$[I + \delta^2 \mathcal{P}_f^j [I - \delta^2 \mathcal{P}_f^i \cdot \mathcal{P}_f^j]^{-1} \cdot \mathcal{P}_f^i] = [I - \delta^2 \mathcal{P}_f^j \cdot \mathcal{P}_f^i]^{-1},$$

it follows that $D_{U_i} J_i$ is given by

$$D_{U_i} J_i = (1 - \delta) \Lambda(i) \cdot [I - \delta^2 \mathcal{P}_f^j \cdot \mathcal{P}_f^i]^{-1} \cdot [I + \delta \mathcal{P}_f^j] \cdot I \quad (18)$$

Next examine $D_{U_i} J_i$ in equation (18). Clearly, since $[I - \delta^2 \mathcal{P}_f^j \cdot \mathcal{P}_f^i]^{-1}$ is invertible it must have full rank of $|S|$. Also, since $[I + \delta \mathcal{P}_f^j]$ has a dominant diagonal, it is invertible and therefore has full rank of $|S|$.

To complete the result, 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. By this fact, $\Lambda(i) \cdot [I + \delta^2 \mathcal{P}_f^j [I - \delta^2 \mathcal{P}_f^i \cdot \mathcal{P}_f^j]^{-1} \cdot \mathcal{P}_f^i]$ has rank of $|R_i|$ while $[I + \delta \mathcal{P}_f^j] \cdot I$ has rank of $|S|$. Applying this fact again to the product of these two establishes that $D_{U_i} J_i(f, U)$ has full rank of $|R_i|$. \blacksquare

Proof of Lemma 2 Let $(f^*, U^*) \in F_C \times \mathbb{R}^{2|S|}$ satisfy $J(f^*, U^*) = 0$. By Lemma 1, the matrix $D_{U_i} J_i(f^*, U^*)$ has full row rank $|R_i|$ for each i . Relying on the fact that row rank and column rank of a matrix coincide, we can find a partition (P_i, Q_i) of C_i for each i such

that $|P_i| = |R_i|$ and $|Q_i| = |C_i| - |R_i|$ with notation $U_P^* = (U_{P_1}^*, U_{P_2}^*)$ and $U_Q^* = (U_{Q_1}^*, U_{Q_2}^*)$ such that

$$\frac{\partial J}{\partial U_P} \text{ has full rank } |R_1| + |R_2| \text{ at } (f^*, U^*).$$

We therefore apply the Implicit Function Theorem to obtain open neighborhoods \mathcal{F} , \mathcal{U}_P and \mathcal{U}_Q of f^* , U_P^* and U_Q^* , respectively, and a locally smooth mapping $\phi : \mathcal{F} \times \mathcal{U}_Q \rightarrow \mathcal{U}_P$ satisfying

$$J(f, (\phi(f, U_Q), U_Q)) = 0$$

where $U_P = \phi(f, U_Q)$ is the unique solution to $J(f, (U_P, U_Q)) = 0$ in \mathcal{U}_P for all $(f, U_Q) \in \mathcal{F} \times \mathcal{U}_Q$. Consequently, the mapping $\Psi : \mathcal{F} \times \mathcal{U}_Q \rightarrow \mathcal{U}_P \times \mathcal{U}_Q$ defined by

$$\Psi(f, U_Q) = (\phi(f, U_Q), U_Q) \tag{19}$$

has the requisite properties. ■

Proof of Lemma 3 Let Y denote the set of f -critical points of Ψ and Z denote the set of f -critical values of Ψ . Then $Z = \Psi(Y)$.

Let ϕ denote the locally smooth mapping that defines Ψ in Equation (19). For each $U_Q \in \mathcal{U}_Q$, the restricted map $\phi(\cdot, U_Q) : \mathcal{F} \rightarrow \mathcal{U}_P$ is also locally smooth. By Sard's theorem, the set of critical values of $\phi(\cdot, U_Q)$ constitutes a subset, $\mathcal{V}_P(U_Q) \subset \mathcal{U}_P$, with Lebesgue measure zero in $\mathbb{R}^{|P_1|+|P_2|}$. By construction,

$$Z = \{(U_P, U_Q) \in \mathcal{U}_P \times \mathcal{U}_Q : U_P \in \mathcal{V}_P(U_Q)\}$$

We proceed to show first that Z is a measurable set in $\mathbb{R}^{2|S|}$, and second, using Fubini's Theorem, that Z has measure 0 given that it is measurable.

Regarding measurability of Z , we follow a lead in the proof of Sard's theorem in Milnor (1965), p. 18. The set $(\mathcal{F} \times \mathcal{U}_Q) \setminus Y$ is open relative to $\mathcal{F} \times \mathcal{U}_Q$ and thus Y is closed relative to $\mathcal{F} \times \mathcal{U}_Q$. Let $\mathbb{Q}^{2|S|}$ denote the set of points in $\mathbb{R}^{2|S|}$ with rational coordinates and let \mathbb{Q}_{++} denote the set of strictly positive rational numbers. For a point $(x, r) \in \mathbb{Q}^{2|S|} \times \mathbb{Q}_{++}$, let $K_r(x)$ denote the closed ball in $\mathbb{R}^{2|S|}$ with center x and radius r . For any $y = (f, U_Q) \in \mathcal{F} \times \mathcal{U}_Q$, there exists $(x, r) \in \mathbb{Q}^{2|S|} \times \mathbb{Q}_{++}$ with $y \in K_r(x) \subset \mathcal{F} \times \mathcal{U}_Q$. Hence the closed set Y in $\mathcal{F} \times \mathcal{U}_Q$ can be covered by a countable family of compact subsets $K_n, n \in \mathbb{N}$, of $\mathcal{F} \times \mathcal{U}_Q$. For each $n \in \mathbb{N}$, $L_n \equiv \Psi(K_n)$ is a compact subset of Z and $Z = \Psi(Y) = \Psi(\bigcup_n K_n) = \bigcup_n \Psi(K_n) = \bigcup_n L_n$. Hence Z is a countable union of compact sets and therefore measurable.

Now set $Z(U_Q) = \{U_P \in \mathcal{U}_P : (U_P, U_Q) \in Z\}$ for each $U_Q \in \mathcal{U}_Q$. An element $U_P \in \mathcal{U}_P$ is a critical value for the mapping $\phi(\cdot, U_Q)$, if there exists a critical point $f \in \mathcal{F}$ of $\phi(\cdot, U_Q)$,

with $U_P = \phi(f, U_Q)$. But

$$\Psi(f; U_Q) = (\phi(f, U_Q), U_Q) = (U_P, U_Q)$$

Hence U_P is a critical value of $\phi(\cdot, U_Q)$ if and only if (U_P, U_Q) is an f -critical value of Ψ . Therefore, $Z(U_Q) = \mathcal{V}_P(U_Q)$. Since, by construction, $\mathcal{V}_P(U_Q)$ is a measure 0 subset of \mathcal{U}_P (in the space $\mathbb{R}^{|P_1|+|P_2|}$), it follows that the section $Z(U_Q)$ of Z has Lebesgue measure zero in $\mathbb{R}^{|P_1|+|P_2|}$. Since this holds true for any choice of U_Q and since Z is measurable, Fubini's Theorem implies that Z is a Lebesgue measure zero subset of $\mathcal{U}_P \times \mathcal{U}_Q$ in the space $\mathbb{R}^{|C_1|+|C_2|}$. \blacksquare

Proof of Lemma 4 Recall the open sets $\mathcal{F}, \mathcal{U}_P, \mathcal{U}_Q$, and the mapping $\Psi : \mathcal{F} \times \mathcal{U}_Q \rightarrow \mathcal{U}_P \times \mathcal{U}_Q$ with the measure zero set Z of f -critical values to which Lemma 4 applies. Suppose that \hat{U} is an f -regular value of Ψ . That is, suppose $\hat{U} \notin Z$. We proceed to show that each \hat{f} with

$$\hat{f} \in \{f \in \mathcal{F} : J(f, \hat{U}) = 0\}$$

is an isolated point in this set.

Because of $\hat{U} \notin Z$, (\hat{f}, \hat{U}_Q) is an f -regular point of Ψ , i.e. $D_f \Psi$ has full rank of $|R_1| + |R_2|$ at (\hat{f}, \hat{U}_Q) . Hence the derivative $D_f \phi(\cdot, \hat{U}_Q)$ has full rank at \hat{f} . Therefore, by the Inverse Function Theorem, there exist open neighborhoods $\hat{\mathcal{F}}$ and $\hat{\mathcal{U}}_P$ with $\hat{f} \in \hat{\mathcal{F}} \subseteq \mathcal{F}$ and $\hat{U}_P \in \hat{\mathcal{U}}_P \subseteq \mathcal{U}_P$, respectively, such that $\hat{\phi}(\cdot, \hat{U}_Q)$, the restriction of $\phi(\cdot, \hat{U}_Q)$ to $\hat{\mathcal{F}}$, is a diffeomorphism from $\hat{\mathcal{F}}$ onto $\hat{\mathcal{U}}_P$.

We now claim that \hat{f} is the only point in the set

$$\hat{f} \in \{f \in \hat{\mathcal{F}} : J(f, \hat{U}) = 0\}$$

For suppose otherwise. Suppose that g also lies in this set. Then $g, \hat{f} \in \hat{\mathcal{F}}$ and $J(\hat{f}, \hat{U}) = J(g, \hat{U}) = 0$ where $\hat{U}_P \in \hat{\mathcal{U}}_P$. Therefore $\hat{U}_P = \hat{\phi}(\hat{f}, \hat{U}_Q) = \hat{\phi}(g, \hat{U}_Q)$. Since $\hat{\phi}(\cdot, \hat{U}_Q)$ is a diffeomorphism from $\hat{\mathcal{F}}$ onto $\hat{\mathcal{U}}_P$, it follows that $g = \hat{f}$ as asserted. \blacksquare

Proof of Lemma 5 Fix $U \in \mathcal{W}$. Suppose that $\mathcal{X}(U) \cap F_C$ is infinite. Then there exists a sequence $\{f_k\}_{k=1}^{\infty}$ of pairwise distinct elements of $\mathcal{X}(U) \cap F_C$. Now $\mathcal{X}(U) \cap F_C$ is contained in the closure of F_C , a compact set. Passing to subsequences if necessary it follows that $f_k \rightarrow f$ with f belonging to the closure of F_C . Hence $f \in G_C$. Since $J(f_k, U) = 0$ for all k and since J is smooth, $J(f, U) = 0$. Therefore, $f \in \mathcal{X}(U)$ and there exists an open set \mathcal{F}^n with $\mathcal{F}^n \cap \mathcal{X}(U) = \{f\}$. But since each f_k belongs to $\mathcal{X}(U)$, $f_k \rightarrow f$ and \mathcal{F}^n is

open, it must also be the case that $f_k \in \mathcal{F}^n \cap \mathcal{X}(U)$ for infinitely many k . However, this contradicts the fact that $\mathcal{F}^n \cap \mathcal{X}(U) = \{f\}$. Hence to the contrary, $\mathcal{X}(U) \cap F_C$ has to be finite. ■

Proof of Lemma 6 Fix U and let f° be an MPE with support C . Choose the normalization $D = (D_1, D_2)$ and $R_i = C_i \setminus D_i$ for each i as before. Recall that J is defined by equation (14). In particular, $f_i^\circ(s_i|s_j) > 0$ for each i and each $s = (s_i, s_j) \in R_i$. By (11),

$$J_{i s}(f^\circ, U) = \frac{\partial V_i(f|s_j)}{\partial f_i(s_i|s_j)} \Big|_{f=f^\circ} = 0. \quad (20)$$

Since (20) must be satisfied for each $i = 1, 2$ and each $s \in R_i$, it follows that $E_C(U) \subseteq \mathcal{X}(U) \cap E_C$. ■

References

- [1] Bhaskar, V. and F. Vega-Redondo (2002), “Asynchronous Choice and Markov Equilibria” *Journal of Economic Theory*, 103.
- [2] Dutta, P. (1995), “A Folk Theorem for Stochastic Games,” *Journal of Economic Theory*, 66, 1-32.
- [3] Friedman, J.W. (1986), *Game Theory with Applications to Economics*. Oxford University Press: New York and Oxford.
- [4] Fudenberg, D., and J. Tirole (1991), *Game Theory*. The MIT Press: Cambridge, MA, and London.
- [5] Govindan, S. and A. McLennan (1999), “Generic Finiteness of Equilibrium Outcome Distributions,” for Two Person Game Forms with Zero Sum and Common Interest Utilities,” mimeo.
- [6] Haller, H. and R. Lagunoff (2000), “Genericity and Markovian Behavior in Stochastic Games,” *Econometrica*, 68, 1231-48.
- [7] Harsanyi, J.C. (1973), “The Oddness of Equilibrium Points: A New Proof,” *International Journal of Game Theory*, 2, 235-250.
- [8] Herings, P. J.-J. and R.J.A.P. Peeters (2004), “Stationary Equilibria in Stochastic Games: Structure, Selection, and Computation,” *Journal of Economic Theory*, 118, 32-60.

- [9] Lagunoff, R. and A. Matsui (1997), “Asynchronous Choice in Repeated Coordination Games,” *Econometrica*, 65: 1467-78.
- [10] Mas-Colell, A. (1985), *The Theory of General Economic Equilibrium: A Differentiable Approach*. Cambridge University Press: Cambridge et al.
- [11] Maskin, E., and J. Tirole (1987), “A Theory of Dynamic Oligopoly, Part III: Cournot Competition”, *European Economic Review* 31, 947-968.
- [12] Maskin, E. and J. Tirole (1988a), “A Theory of Dynamic Oligopoly, I: Overview and Quantity Competition with Large Fixed Costs,” *Econometrica*, 56, 549-70.
- [13] Maskin, E. and J. Tirole (1988b), “A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand Curves and Fixed Costs, ” *Econometrica*, 56, 571- 600.
- [14] Milnor, J.W. (1965), *Topology From the Differentiable Point of View*. The University Press of Virginia: Charlottesville.
- [15] Park, I.-U. (1997), “Generic Finiteness of Equilibrium Outcome Distributions for Sender-Receiver Cheap-Talk Games,” *Journal of Economic Theory*, 76, 431-448.
- [16] 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.
- [17] Wilson, R. (1971), “Computing Equilibria in N -Person Games,” *SIAM Journal of Applied Mathematics*, 21, 80-87.
- [18] Yoon, K. (2001), “A Folk Theorem for Asynchronously Repeated Games,” *Econometrica*, 69: 191-200.