

Markov Perfect Equilibria in Repeated Asynchronous Choice Games

Hans Haller* and Roger Lagunoff†

April 16, 1997‡

Abstract

This paper examines the issue of multiplicity of equilibria in alternating move repeated games with two players. Such games are canonical models of environments with repeated, asynchronous choices due to inertia or replacement. We focus our attention on *Markov Perfect equilibria (MPE)*. These are Perfect equilibria in which individuals condition their actions on payoff-relevant state variables. 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. We also compare the MPE to non-Markovian equilibria and to the (trivial) MPE of standard repeated games. Unlike the latter, it is often true when moves are asynchronous that Pareto inferior stage game equilibrium payoffs cannot be supported in MPE. Also, MPE can be constructed to support cooperation in a Prisoner’s Dilemma despite limited possibilities for constructing punishments.

Journal of Economic Literature Classification Numbers: C72, C73.

Keywords: repeated games, asynchronous choice, turn-taking games, 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, and especially Akihiko Matsui for helpful conversations.

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

Apart from the standard model of repeated play, there are many 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. Another example of asynchronous play are the birth/death processes in evolutionary models. In such models, replacement occurs at independent, random times forcing decisions at the time of entry to be asynchronous.

For whatever reason that asynchronous play exists, its effect in repeated situations is sometimes different than in the standard model. Lagunoff and Matsui (1997) show that in repeated games with asynchronous choices, if the stage game is one of pure coordination, i.e., a game in which the payoffs to all players in each action profile are identical, then an “Anti-Folk Theorem” obtains. That is, there is a unique perfect equilibrium payoff if players are sufficiently patient. However, more generally Dutta (1995) proves a Folk Theorem for stochastic games when a full dimensionality condition holds.² Since sequential move games have stochastic game representations, Dutta’s result shows that asynchronous timing does not, by itself, typically reduce the multiplicity of equilibria when players are very patient.

The present paper re-examines the multiplicity issue in repeated games with asynchronous choices. We focus our attention on those equilibria which arise when players’ strategies are Markovian in the sense that, rather than conditioning their behavior on the entire observed history, players condition only on payoff-relevant states.³ We examine the timing structure considered by Maskin and Tirole (1987, 1988a,b) as a canonical model of asynchronous repetition. Two players take turns moving in a stationary environment. In turn-taking games, the natural state variable is the move of other player(s) in the previous period. In using Markovian strategies, the players need not coordinate their behavior to depend on the past. Indeed, they need not have observed the past at all. Each player merely needs to know “where he is” in the game. This requires only that he understand what are his feasible actions at any given time and what are their consequences.

¹See, for example, Abreu (1988), Fudenberg and Maskin (1986), and Aumann (1981).

²In *Stochastic games* the stage game evolves according to a transition function which maps (probabilistically) the current state and action profile into next period’s state.

³See Fudenberg and Maskin (1995) for an alternative notion of Markovian behavior.

Our main result is that the number of *Markov Perfect Equilibria* in this game — with the abbreviation MPE used in the sequel — is generically finite with respect to stage game payoffs. What is unusual about this result is the limited degree of freedom offered by stage game payoffs for genericity arguments involving dynamic equilibria. As we observed, the alternating move repeated game is a stochastic game. Hence, a more natural candidate for a genericity argument is the larger dimensional space of state-dependent payoffs. The problem is that the alternating move game is “non-generic” in this larger space. In fact, the largest degree of freedom is precisely given by the space of stage game payoffs, and so we are restricted to it for proving the result.

The result contrasts with models of non-Markovian strategies, including those with low complexity. Kalai and Stanford (1988) show, for example, 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. Of course, Markovian strategies are themselves simple automata. The difference is that the restriction to Markovian strategies rules out the possibility to extend memory, by using states of the automaton, beyond what is encoded in the states of the game. Since even one period of past history can support a continuum of equilibria in some games, our result suggests a striking discontinuity between Markovian and non-Markovian behavior.

Naturally, in the standard model of repeated play, if the players do not use past history then only Nash equilibria of the stage game can arise. In that case, since the only Markovian strategies are stage game strategies, only Nash equilibria of the stage game are Markov Perfect in the sense which we described. In such a case, there is a well known result of Wilson (1970) and Harsanyi (1971) who show that the number of Nash equilibria in any normal form (stage game) is generically finite with respect to payoffs. Hence, in a trivial sense, Markovian equilibria of the standard model of a repeated game are also generically finite, and so the discontinuity described above exists in the standard model as well.

Yet, the fact that nontrivial state variables exist in the asynchronous choice case but not in the standard model does make a significant difference. We show that despite the relatively small number of equilibria, Markovian behavior can support cooperation in some Prisoner’s Dilemma games, whereas it rules out Pareto inferior Nash equilibria in some coordination games. Therefore, while the restriction to Markov behavior pares down the equilibrium set in both models, states variables in the asynchronous model admit enough flexibility to support desirable outcomes which could not arise in the stage game, and eliminate some undesirable outcomes of the stage game which standard repetition cannot prevent.⁴

⁴A recent result by Dutta (1996) establishes efficiency of MPE in stochastic games. His result utilizes a condition called “state symmetry” in which all individuals essentially rank the states in the same way (in their preferences) under optimal behavior. Examples include common pool resource extraction problems.

This paper is organized as follows. Section 2 introduces definitions and notation and describes existing results which are relevant for comparison. Section 3 demonstrates differences between Markovian and non-Markovian equilibria and between synchronous and asynchronous timing. Three results for 2×2 stage games are described. The first demonstrates why the equilibrium set is generically finite in this special case. The second examines coordination games and shows that sometimes, Pareto inferior Nash equilibria of the stage game cannot result in MPE. The third shows that sometimes MPE can be constructed to support cooperation in a Prisoner's Dilemma despite the limited possibility for punishments. Section 4 examines the logic of the main result, provides comments on concepts of genericity, and finally gives the proof of the main result for general, two player alternating move games.

2 The Canonical Model

2.1 Preliminaries

Consider the alternating move game of which stage payoff is given by $G = \langle S_1, S_2, u_1, u_2 \rangle$ where S_i , $i = 1, 2$, is a finite set of actions of player i , and, letting $S = S_1 \times S_2$, $u_i : S \rightarrow \mathfrak{R}$ is the utility function of player i . Alternatively, let \mathcal{U}_i denote i 's $S \times 1$ utility vector so that $\mathcal{U}_{is} = u_i(s)$. Then $\mathcal{U} = (\mathcal{U}_1, \mathcal{U}_2)$ is an element of \mathfrak{R}^{2S} . Abusing notation a bit, we will use S_i to refer to both the set and the cardinality of the set of i 's actions in the stage game.

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.

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 $\sum_{t=0}^{\infty} (1 - \delta) \delta^t u_i(s^t)$ given common discount factor δ . Let H_1 denote the set of all histories ending in odd numbered periods so that $h \in H_1$ implies that for some t , $h = (s^0, s^1, s^2, \dots, s^{2t+1})$. Similarly, H_2 is the set of all histories ending in even numbered periods so that $h \in H_2$ implies $h = (s^0, s^1, s^2, \dots, s^{2t})$. 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

Our results differ, however, since the alternating move Prisoner's Dilemma is apparently not state symmetric.

current history is h .

To define the payoffs, we need one more useful bit of notation. Let ι denote an ‘‘indicator’’ of the current mover so that $\iota(t) = 1$ if t is odd, and $\iota(t) = 2$ if t is even. Given $f = (f_1, f_2)$, player i 's payoff after history h^t is given by

$$V_i(f|h^t) \equiv (1 - \delta)E \left[\sum_{\tau=0}^{\infty} \delta^{t+\tau} \sum_{s_{\iota(t+\tau)}^{t+\tau}} f_{\iota(t+\tau)}(s_{\iota(t+\tau)}|h^{t+\tau}) u_i(s) \right] \quad (1)$$

where the expectation is taken with respect to the distribution induced by the f .

2.2 Perfect Equilibrium and The Dutta Folk Theorem

A *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 $h \in H$, i.e.,

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

for any of player i 's strategies f_i . Now let

$$m_i(h^t) = \lim_{\delta \rightarrow 1} \inf_{f_j} \sup_{f_i} V_i^\delta(f|h^t)$$

where V_i^δ denotes explicit dependence on discount factor δ . The payoff $m_i(h^t)$ denotes the long run minimax payoff for player i given history h^t . Let W denote the convex hull of the stage game payoff vectors, and define the set of feasible payoffs given h^t that strictly dominate the long run minimax payoff by

$$F(h^t) = \left\{ w \in W \mid w_i > m_i(h^t), i = 1, 2 \right\}$$

A result of Dutta (1995) proves a Folk Theorem for general stochastic games. In our model, he shows that any payoff in $F(h^t)$ can be arbitrarily approximated by perfect equilibrium payoffs for sufficiently patient players if the feasible payoffs and the minimax payoffs do not vary across histories, and if the game satisfies a full dimensionality requirement. More precisely,

Theorem (Dutta (1995)) *Suppose that for each $i = 1, 2$, and all histories h^t , $m_i(h^t) = m_i$ and $F(h^t) = F$ with $\dim(F) = 2$. For any $w \in F$ and any $\epsilon > 0$ there is a $\bar{\delta}$ such that if $\delta \geq \bar{\delta}$, then there exists a perfect equilibrium f with $V_i(f|h^t, \delta) > w_i - \epsilon$, $\forall h^t \in H$.*

2.3 Markov Perfect Equilibria

The application of Dutta’s result here generically resolves the multiplicity issue when all perfect equilibria of the turn-taking game 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. 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,

$$V_1(f|s_i) = \sum_{s_j} f_j(s_j|s_i) [(1 - \delta)u_1(s_i, s_j) + \delta V_1(f|s_j)] \quad (2)$$

A similar expression for $i = 2$ is obtained by switching subscripts. A *Markov Perfect Equilibrium* is a perfect equilibrium in which the strategies are Markovian.

Existence of Markov Perfect equilibria is a standard result.⁵ The main result of this paper is:

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

Notice 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.

As for the proof, which is contained in Section 4, there are two difficulties which prevent a straightforward application of Harsanyi’s proof to Markovian equilibria in repeated games with asynchronous choice. First, 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. Second and more importantly, the “genericity” in our result refers to the dimension of payoffs in the stage game rather than in the stochastic game. To see the difference, consider the stage game in Figure 1 below.

⁵See, for example, Friedman (1986), Fudenberg and Tirole (1991), or Lagunoff and Matsui (1997).

		II	
		s_2	s'_2
I	s_1	5, 3	20, 2
	s'_1	10, 1	0, 4

Figure 1

The stochastic game representation of the 2×2 game in Figure 1 is seen below in Figure 2. The state space is the set $S_1 \cup S_2$. If the current state is s_2 , for example, then this means that individual 2 moved in the previous period and took action s_2 . The arrows denote the transition laws.




Figure 2

If one were to prove that finite state stochastic games have generically finitely many Markovian equilibria, then the term “generically” would reasonably apply to the state dependent utilities. In this case, we would have to consider all possible payoff vectors in \mathbb{R}^{32} since there are 8 possible stage game payoffs, and 4 possible states. However, the redundancy in payoffs is clear from Figure 2; the alternating move game gives not more than 8 degrees of freedom. That is to say, the stochastic game representation of repeated alternating moves will not be a “generic” stochastic game! Hence, we are restricted to consider variations in the original stage game only.⁶

⁶Even if one allows the set of feasible actions to vary across states, there will still be some redundancy in payoffs.

3 Comparison with Standard Repeated Games

		2	
		C	D
1	C	a_1, a_2	b_1, d_2
	D	d_1, b_2	c_1, c_2

Figure 3

In this Section we give some intuition for the result in the 2×2 stage game in Figure 3 above. We also demonstrate how MPE of this model differ from: (1) Perfect equilibria of the model, and (2) MPE of the standard repeated game. For the first comparison, observe that if the stage game above satisfies Dutta’s dimensionality condition (which is much weaker than having the payoffs spanning \mathfrak{R}^8), then any individually rational payoff may be shown to be supported by a Perfect equilibrium of the turn-taking game if players are patient enough. For the second comparison, note that Nash equilibria of the stage game are (trivial) Markov Perfect equilibria of the standard repeated game model.

Now combine the restriction to Markovian behavior with a turn-taking repetition of the stage game. 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 . Clearly, there are only finitely many “pure strategy” configurations, i.e., finitely many cases in which $\alpha_i \in \{0, 1\}$ and $\beta_i \in \{0, 1\}$. However, this is almost always true for mixed strategies as well.

Claim 1 *Except on a set of payoffs (a_i, b_i, c_i, d_i) with measure zero in \mathfrak{R}^4 , there are no MPE in which player i uses a strictly mixed strategy, i.e., there are no MPE in which $0 < \alpha_i < 1$ and $0 < \beta_i < 1$.*

To see this, 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 (2), observe that

$$V_{iC} = \alpha_i ((1 - \delta)a_i + \delta W_{iC}) + (1 - \alpha_i) ((1 - \delta)d_i + \delta W_{iD}) \tag{3}$$

and

$$V_{iD} = \beta_i ((1 - \delta)c_i + \delta W_{iD}) + (1 - \beta_i) ((1 - \delta)b_i + \delta W_{iC}) \tag{4}$$

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

$$V_{iC} = (1 - \delta)a_i + \delta W_{iC} = (1 - \delta)d_i + \delta W_{iD}, \tag{5}$$

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

$$V_{iD} = (1 - \delta)c_i + \delta W_{iD} = (1 - \delta)b_i + \delta W_{iC}. \quad (6)$$

However, equations (5) and (6) are mutually consistent only if $d_i + b_i = a_i + c_i$. The 4-tuple of payoffs for which $d_i + b_i = a_i + c_i$ spans a 3-dimensional hyperplane in \mathfrak{R}^4 .

It can also be shown that the incompletely mixed MPE are finite in number except on null sets in \mathfrak{R}^8 . Observe that nothing in the above argument relies on the specific structure of the stage game. The linearity of the system shown in equations (5) and (6) allows us to simply compare equations and unknowns. In the general case, we lose the linearity in strategies but retain linearity in stage payoffs. The following two claims make other useful comparisons to the standard synchronous model of repetition.

Claim 2 *Suppose that the game is a coordination game, i.e., $a_i > d_i$ and $c_i > b_i$ for each i . Suppose further that $a_i > c_i$ so that (C, C) is the Pareto superior Nash equilibrium. Then, if $b_i > d_i$ for each i , and if players are sufficiently patient, then the Nash equilibrium (D, D) cannot comprise any MPE of the repeated turn-taking game.*

The argument for this claim goes as follows. Suppose, by contradiction, that (D, D) is supported by a MPE. First, we must have $\beta_i = f_i(D|D) = 1$ for both $i = 1, 2$. Then $V_{iD} = c_i$ and

$$V_{iC} = \alpha_i[(1 - \delta)a_i + \delta W_{iC}] + (1 - \alpha_i)[(1 - \delta)d_i + \delta c_i], \text{ and} \quad (7)$$

$$W_{iC} = \alpha_j[(1 - \delta)a_i + \delta V_{iC}] + (1 - \alpha_j)[(1 - \delta)b_i + \delta c_i]. \quad (8)$$

Also, the incentive constraint to choose D requires

$$c_i \geq (1 - \delta)b_i + \delta W_{iC}. \quad (9)$$

Next notice that $\alpha_i \neq 1$ must hold for both i . For otherwise, some player always has a best response to remain in (C, C) by choosing $\alpha_i = 1$ if the other player does so. But then (D, D) cannot be reached from (C, C) which is a contradiction. If $1 > \alpha_i > 0$ for some i then it follows from the equilibrium condition that

$$(1 - \delta)a_i + \delta W_{iC} = (1 - \delta)d_i + \delta c_i$$

which, combined with (9) gives $a_i + c_i \leq d_i + b_i$ which is clearly a contradiction. Finally, suppose that $\alpha_i = 0$ for both i . Then

$$V_{iC} = (1 - \delta)d_i + \delta c_i, \text{ and} \quad (10)$$

$$W_{iC} = (1 - \delta)b_i + \delta c_i \quad (11)$$

But since $b_i > d_i$, it follows that $W_{iC} > V_{iC}$. However, since $V_{iC} < (1 - \delta)a_i + \delta W_{iC}$, player i could achieve a strictly higher payoff by choosing $\alpha_i = 1$ than by choosing $\alpha_i = 0$. Hence, $\alpha_i = 0$, $i = 1, 2$ cannot comprise a MPE. As we have exhausted all cases in which (D, D) might possibly comprise a MPE, it cannot occur.

Claim 3 *Suppose that the game is a Prisoner's Dilemma game, i.e., $d_i > a_i > c_i > b_i$ for each i . If $a_i + c_i > b_i + d_i$ for each i , then for sufficiently patient players the profile (C, C) comprises some MPE.*

To prove Claim 3, we construct 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 so that each player chooses a mixed strategy when the other chose D last period. 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) :

$$a_i > (1 - \delta)d_i + \delta W_{iD} \tag{12}$$

$$V_{iD} = (1 - \delta)c_i + \delta W_{iD} = (1 - \delta)b_i + \delta a \tag{13}$$

$$W_{iD} = \beta_j[(1 - \delta)c_i + \delta V_{iD}] + (1 - \beta_j)[(1 - \delta)d_i + \delta a] \tag{14}$$

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.

4 Main Result: the Logic and the Proof

4.1 Markov Transitions

To take advantage of the 2×2 intuition, an alternative to the expression (2) 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:

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

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 II chooses D after I chose D , and player I chooses C after II 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

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

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:

$$V_i(f|s_j) = \pi_s \cdot \left[\sum_{t=0}^{\infty} (1 - \delta) \delta^{2t} [\mathcal{P}_f^i \cdot \mathcal{P}_f^j]^t \cdot \mathcal{P}_f^i \cdot [I + \delta \mathcal{P}_f^j] \right] \cdot \mathcal{U}_i \quad (15)$$

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

$$\begin{aligned}
A_f^i &= \left[\sum_{t=0}^{\infty} (1-\delta)\delta^{2t} [\mathcal{P}_f^i \cdot \mathcal{P}_f^j]^t \cdot \mathcal{P}_f^i \cdot [I + \delta\mathcal{P}_f^j] \right] \\
&= (1-\delta)[I - \delta^2(\mathcal{P}_f^i \cdot \mathcal{P}_f^j)]^{-1} \cdot \mathcal{P}_f^i \cdot [I + \delta\mathcal{P}_f^j]
\end{aligned} \tag{16}$$

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 (15) can now be rewritten as

$$V_i(f|s_j) = \pi_s \cdot A_f^i \cdot \mathcal{U}_i. \tag{17}$$

From (17) it is clear that although the dynamic payoff is not multilinear in strategy profiles, f , it is linear in stage game payoffs \mathcal{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 (17) 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.

4.2 Genericity

For the result, we adopt 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 . There are two notions of genericity: one is measure-theoretic, and the other is topological. The standard notion of genericity in the measure-theoretical sense assumes that B is a measurable subset of a Euclidean space \mathfrak{R}^ℓ with $\lambda^\ell(B) > 0$ where λ^ℓ is Lebesgue measure on \mathfrak{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.

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 may be responded to in part by pragmatism. 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.⁷ We therefore opt for the measure-theoretic notion with Lebesgue measure as the canonical measure, in the tradition of Harsanyi (1971), Wilson (1971), and others. Whether the Theorem would go through as well when the topological notion were used, is an open question. For more on genericity notions, see Mas-Colell (1985).

4.3 The Proof

Recall from equation (17) that the payoff for $i = 1, 2$, is given by $V_i(f|s_j) = \pi_s \cdot A_f^i \cdot \mathcal{U}_i$ where A_f^i is described in equation (16). Obviously, the right-hand side of (17) is C^∞ in the pair (f, \mathcal{U}_i) . Next let us exhibit some of the first order conditions for maximization of (17).

Lemma 1 *A necessary condition for f^* to be an MPE is that*

$$\left. \frac{\partial V_i(f^*|s_j)}{\partial f_i(s_i|s_j)} \right|_{\sum f_i(t_i|s_j)=1} = 0, \quad (18)$$

holds for each (s_i, s_j) and each i for which $0 < f_i^(s_i|s_j) < 1$.*

Proof Clear by construction of V_i . □

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

Until further notice, we shall consider several cases distinguished by carriers. So let us fix some carrier $C = (C_1, C_2)$ and denote by $F(C) = F_1(C_1) \times F_2(C_2)$ the set of $f \in F$ with carrier C : $f \in F(C)$ if and only if $C = C(f)$.

⁷For 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. A set exhibiting these features can be obtained through a Cantor-like construction: Cf. Example 4.5 in Gilles, et al. (1997).

Moreover, for each i and $s_j \in S_j$, we fix some action $d(s_j) \in S_i$ for player i satisfying $(d(s_j), s_j) \in C_i$. Let $D_i = \{(d(s_j), s_j) : s_j \in S_j\}$ denote the set of “normalizing” profiles for player i . Further set $R_i = C_i \setminus D_i$. Notice then that the case $R_1 = \emptyset$ and $R_2 = \emptyset$ implies that f_i is a pure strategy for each i . For any utility vector \mathcal{U} , the game has a finite number of equilibria in pure strategies, since there are only finitely many pure strategy profiles. The special case $R_i \neq \emptyset$, $R_j = \emptyset$ with $i \neq j$ can be dealt with in a way analogous to the treatment of the case where both sets are non-empty. So let us proceed under the assumption that $R_1 \neq \emptyset$, $R_2 \neq \emptyset$.

Observe that there are at most $2S$ equations of the form (18) above. We now show that (18) can be replaced by a system of algebraic equations that have to hold for any MPE. Varying $s_j \in S_j$ for $j \neq i$, $i = 1, 2$, we pin down $S_1 + S_2$ normalizing equations satisfying

$$f_i(d(s_j)|s_j) = 1 - \sum_{s_i \neq d(s_j)} f_i(s_i|s_j), \quad d(s_j) \in S_i, \quad s_j \in S_j, \quad j \neq i, \quad i = 1, 2. \quad (19)$$

We can then express the remaining $2S - S_1 - S_2$ equations as

$$f_i(s_i|s_j) f_i(d(s_j)|s_j) (\pi_s \cdot B_f(s_i|s_j) \cdot \mathcal{U}_i) = 0 \quad (20)$$

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

$$\begin{aligned} B_f(s_i|s_j) &= \frac{\partial}{\partial f_i(s_i|s_j)} [A_f^i] \\ &= \frac{\partial}{\partial f_i(s_i|s_j)} \left[(1 - \delta) [I - \delta^2 \mathcal{P}_f^i \cdot \mathcal{P}_f^j]^{-1} \cdot \mathcal{P}_f^i \cdot [I + \delta \mathcal{P}_f^j] \right] \\ &= (1 - \delta) \left([I - \delta^2 \mathcal{P}_f^i \cdot \mathcal{P}_f^j]^{-1} \cdot \left[\frac{\partial \mathcal{P}_f^i}{\partial f_i(s_i|s_j)} \right] \cdot [I + \delta \mathcal{P}_f^j] \right. \\ &\quad \left. - [I - \delta^2 \mathcal{P}_f^i \cdot \mathcal{P}_f^j]^{-1} \right. \\ &\quad \left. \cdot \left[-\delta^2 \frac{\partial \mathcal{P}_f^i}{\partial f_i(s_i|s_j)} \cdot \mathcal{P}_f^j \right] \cdot [I - \delta^2 \mathcal{P}_f^i \cdot \mathcal{P}_f^j]^{-1} \cdot \mathcal{P}_f^i \cdot [I + \delta \mathcal{P}_f^j] \right) \\ &= (1 - \delta) [I - \delta^2 \mathcal{P}_f^i \cdot \mathcal{P}_f^j]^{-1} \cdot \left[\frac{\partial \mathcal{P}_f^i}{\partial f_i(s_i|s_j)} \right] \\ &\quad \cdot [I + \delta^2 \mathcal{P}_f^i \cdot [I - \delta^2 \mathcal{P}_f^i \cdot \mathcal{P}_f^j]^{-1} \cdot \mathcal{P}_f^i] \cdot [I + \delta \mathcal{P}_f^j]. \end{aligned} \quad (21)$$

Since $f_i(s_i|s_j) > 0$ for every $(s_i, s_j) \in D_i$, then for any $(s_i, s_j) \notin D_i$, either $f_i(s_i|s_j) = 0$ or $0 < f_i(s_i|s_j) < 1$. Hence by (18), all of the $2S - S_1 - S_2$ equations (20) have to be satisfied at an MPE. Combined with the $S_1 + S_2$ normalizing equations (19), a total of $2S$ equations

results.

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 $f_i(s_i|s_j)$. It turns out that $\frac{\partial \mathcal{P}_f^i}{\partial f_i(s_i|s_j)}$ is a $S \times S$ matrix that assigns a 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 a -1 to entries of column $(d(s_j), s_j)$ which are adjacent to profile s (which are also, by construction, adjacent to profile $(d(s_j), s_j)$). Zeros 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 5

After fixing carriers and normalizations, only variation in the R_i -components of f_i matters. Let $F^* = F_1^* \times F_2^*$ be an open subset of $\mathfrak{R}^{R_1+R_2} = \mathfrak{R}^{R_1} \times \mathfrak{R}^{R_2}$ such that F^* contains the corresponding lower-dimensional joint strategy space and all f -related terms in (15) — (21) are well defined. Furthermore, let \mathfrak{R}^{2S} be decomposed into $\mathfrak{R}^{2S} = \mathfrak{R}^* \times \mathfrak{R}^{**}$ where

$$\mathfrak{R}^* = \mathfrak{R}^{R_1} \times \mathfrak{R}^{R_2} \text{ with elements } \mathcal{U}^* = (\mathcal{U}_1^*, \mathcal{U}_2^*) \text{ where } \mathcal{U}_i^* = (\mathcal{U}_{is})_{s \in R_i} \text{ for } i = 1, 2 \text{ and}$$

$$\mathfrak{R}^{**} = \mathfrak{R}^{S \setminus R_1} \times \mathfrak{R}^{S \setminus R_2} \text{ with elements } \mathcal{U}^{**} = (\mathcal{U}_1^{**}, \mathcal{U}_2^{**}) \text{ where } \mathcal{U}_i^{**} = (\mathcal{U}_{is})_{s \notin R_i} \text{ for } i = 1, 2.$$

Now define the mapping $M : F^* \times \mathfrak{R}^{2S} \rightarrow \mathfrak{R}^*$ by $M(f, \mathcal{U}) = (M_{is}(f, \mathcal{U}))_{s=(s_i, s_j) \in R_i; i=1,2}$ where

$$M_{is}(f, \mathcal{U}) = f_i(s_i|s_j) f_i(d(s_j)|s_j) (\pi_s \cdot B_f(s_i|s_j) \cdot \mathcal{U}_i) \quad (22)$$

equals the expression in (20).

Lemma 2 For any given utility vector $\mathcal{U} = (\mathcal{U}_1, \mathcal{U}_2)$, let $\Gamma_{\mathcal{U}, C}$ be the set of solutions $f \in F^*$ to the equations (20), with $s \in R_i; i = 1, 2$. Then:

1. $f \in \Gamma_{\mathcal{U}, C}$ if and only if $M(f, \mathcal{U}) = 0$;
2. if f represents an MPE with carrier C , then $f \in \Gamma_{\mathcal{U}, C}$.

Proof Statement 1 is an immediate consequence of the definition of $\Gamma_{\mathcal{U}, C}$ and M . The argument for statement 2 has been provided immediately after equation (21). \square

Clearly, M is a smooth mapping. Next, let $D_{\mathcal{U}}M(f, \mathcal{U})$ denote the Jacobian matrix with respect to \mathcal{U} . This Jacobian has the form:

$$D_{\mathcal{U}}M(f, \mathcal{U}) = \begin{bmatrix} J_{\mathcal{U}}[11] & 0 \\ 0 & J_{\mathcal{U}}[22] \end{bmatrix}$$

Figure 6

Here, $J_{\mathcal{U}}[ij]$ denotes the $R_i \times S$ block of partials

$$\frac{\partial M_{is}(f, \mathcal{U})}{\partial \mathcal{U}_{js'}} = \frac{\partial [f_i(s_i | s_j) f_i(d(s_j) | s_j) \pi_s \cdot B_f(s_i | s_j) \cdot \mathcal{U}_i]}{\partial \mathcal{U}_{js'}}, \quad j = 1, 2, \quad s' \in S, \quad s \in R_i \quad (23)$$

Observe that if $j \neq i$, then the block $J_{\mathcal{U}}[ij]$ has zero entries everywhere. Notice that each entry in $J_{\mathcal{U}}[ii]$ assumes the form

$$\frac{\partial M_{is}(f, \mathcal{U})}{\partial \mathcal{U}_{is'}} = f_i(s_i | s_j) f_i(d(s_j) | s_j) \pi_s \cdot B_f(s_i | s_j) \cdot \mathcal{U}_i \setminus 1_{s'} \quad (24)$$

where $s = (s_i, s_j) \in R_i$ and $\mathcal{U}_i \setminus 1_{s'} = (\mathcal{U}_{is''}, 1_{s'})_{s'' \neq s'}$ is obtained from the utility vector \mathcal{U}_i by setting the s' -th component equal to 1.

Lemma 3 For each i , there exists $W_i \subseteq \mathfrak{R}^S$ with the following properties:

1. W_i is open and its complement has measure zero;

2. if a pair (f, \mathcal{U}) satisfies $\mathcal{U}_i \in W_i$,
then $J_{\mathcal{U}}[ii]$ has full row rank R_i .

Proof Recall that entries of $J_{\mathcal{U}}[ii]$ are given by equation (24). Since for $s \in R_i$, $f_i(s_i|s_j)f_i(c(s_j)|s_j) > 0$ we temporarily drop this constant in the characterization of $J_{\mathcal{U}}[ii]$. 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 5 one can verify that

$$\Lambda_s(i) = \left(0, \dots, 0, \sum_{s' \in \Theta_i(s)} a_{s's}, 0, \dots, 0, - \sum_{s' \in \Theta_i(s)} a_{s's}, 0, \dots, 0 \right) \quad (25)$$

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, $\sum_{s' \in \Theta_i(s)} a_{s's} \neq 0$. Strictly speaking, $f \in F^*$ allows for slightly negative components. However, choosing F^* small enough keeps these deviations sufficiently small so that we can treat \mathcal{P}_f^i and \mathcal{P}_f^j as if they were stochastic matrices.

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 $J_{\mathcal{U}}[ii]$ is given by

$$J_{\mathcal{U}}[ii] = (1 - \delta) \Lambda(i) \cdot [I - \delta^2 \mathcal{P}_f^j \cdot \mathcal{P}_f^i]^{-1} \cdot [I + \delta \mathcal{P}_f^j] \cdot \bar{\mathcal{U}}_i. \quad (26)$$

where, re-indexing the vector \mathcal{U}_i if necessary, $\bar{\mathcal{U}}_i$ is a $S \times S$ matrix $\bar{\mathcal{U}}_i \equiv [\mathcal{U}_i \setminus 1'_s]^T$ with the form

$$\bar{\mathcal{U}}_i = \begin{bmatrix} 1 & \mathcal{U}_{i1} & \mathcal{U}_{i1} & \cdots & \mathcal{U}_{i1} \\ \mathcal{U}_{i2} & 1 & \mathcal{U}_{i2} & \cdots & \mathcal{U}_{i2} \\ \vdots & & \ddots & & \vdots \\ \vdots & & & \ddots & \vdots \\ \mathcal{U}_{iS} & \cdots & \cdots & \mathcal{U}_{iS} & 1 \end{bmatrix}$$

Figure 7

Now observe $J_{\mathcal{U}}[ii]$ in equation (26). 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 . Now define

$$W_i = \{\mathcal{U}_i \in \mathfrak{R}^S : \bar{\mathcal{U}}_i \text{ has full rank of } S\}$$

We now prove the first part of the Lemma which is

Part 1: W_i is an open set of full measure in \mathfrak{R}^S .

To show this, we begin by establishing some necessary conditions for singularity of $\bar{\mathcal{U}}_i$. We then show that these conditions do not hold on an open set of \mathcal{U}_i having full measure in \mathfrak{R}^S .

Write $\bar{\mathcal{U}}_i = (\bar{\mathcal{U}}_{i1}, \dots, \bar{\mathcal{U}}_{iS})$. Note that $\bar{\mathcal{U}}_{is} = \mathcal{U}_i \setminus 1_s$. Suppose $\bar{\mathcal{U}}_i$ is singular. Then there exist $\lambda_1, \dots, \lambda_S \in \mathfrak{R}$ such that $(\lambda_1, \dots, \lambda_S) \neq 0$ and

$$\sum_k \lambda_k \bar{\mathcal{U}}_{ik} = 0. \quad (27)$$

Line by line, (27) amounts to

$$(1 - \mathcal{U}_{ik}) \cdot \lambda_k + \mathcal{U}_{ik} \cdot \sum_j \lambda_j = 0 \quad (28)$$

for $k = 1, \dots, S$. There are two cases to be distinguished.

Case 1. $\sum_j \lambda_j = 0$. Consequently, for all k , (28) reduces to $(1 - \mathcal{U}_{ik})\lambda_k = 0$. Now for some k , $\lambda_k \neq 0$ and, therefore, $\mathcal{U}_{ik} = 1$ or

$\mathcal{U}_i = (\mathcal{U}_{i1}, \dots, \mathcal{U}_{iS}) \in \mathcal{H}_k \equiv \{x \in \mathfrak{R}^S : x_k = 1\}$.

Case 2. $\sum_j \lambda_j \neq 0$. Then (28) rules out $\mathcal{U}_{ik} = 1$. Consider $\mathcal{U}_{ik} \neq 1$. Then: $\lambda_k = 0 \Leftrightarrow \mathcal{U}_{ik} = 0$. We consider two subcases.

Subcase 2A. $\sum_j \lambda_j \neq 0$ and $\mathcal{U}_{ik} = 0$ for all k .
Then the rows of $\bar{\mathcal{U}}_i$ form a basis of \mathfrak{R}^S and $\bar{\mathcal{U}}_i$ is non-singular.

Subcase 2B. $\sum_j \lambda_j \neq 0$ and $\mathcal{U}_{ik} \neq 0$ for some k .
For any such k , $\mathcal{U}_{ik} \neq 0$ and $\lambda_k \neq 0$. This together with (28) implies

$$\lambda_{-k} \neq 0, \quad (29)$$

where $\lambda_{-k} \equiv \sum_{j \neq k} \lambda_j$, and

$$\mathcal{U}_{ik} = -\lambda_k / \lambda_{-k}. \quad (30)$$

But (29) and (30) also hold for $\mathcal{U}_{ik} = 0$.

Now observe that the set of \mathcal{U}_i with $|\bar{\mathcal{U}}_i| = 0$ is, by continuity, a closed subset of \mathfrak{R}^S . Hence, W_i is an open subset of \mathfrak{R}^S .

Also observe that if $\bar{\mathcal{U}}_i$ is singular, either Case 1 or Subcase 2B obtains. In Case 1, \mathcal{U}_i lies in

$$\mathcal{H} = \bigcup_{k=1}^S \mathcal{H}_k,$$

the union of the hyperplanes \mathcal{H}_k of dimension $S - 1$. Hence, \mathcal{H} is a set of measure zero in \mathfrak{R}^S . In Subcase 2B, the right-hand side of (30) is homogeneous of degree zero with respect to the qualifying weights $\lambda = (\lambda_1, \dots, \lambda_S)$. Normalizing $\sum_j \lambda_j = 1$ allows to recover the λ_k from (29) and (30) as follows: for each k , (30) can be rewritten

$$\mathcal{U}_{ik} \cdot (1 - \lambda_k) = -\lambda_k.$$

Under the conditions of Subcase 2B, this equation can be solved for λ_k , yielding

$$\lambda_k = -\frac{\mathcal{U}_{ik}}{1 - \mathcal{U}_{ik}}. \quad (31)$$

Combining (30) and (31) implies

$$\mathcal{U}_{iS} = -\frac{1 - \lambda_{-S}}{\lambda_{-S}} = \frac{1 + \sum_{k \neq S} \left(\frac{\mathcal{U}_{ik}}{1 - \mathcal{U}_{ik}} \right)}{\sum_{k \neq S} \frac{\mathcal{U}_{ik}}{1 - \mathcal{U}_{ik}}}$$

Hence Subcase 2B is contained in the graph \mathcal{G} of a measurable function whose domain is a measurable subset of \mathfrak{R}^{S-1} and whose range is \mathfrak{R} . By Fubini's Theorem, such a graph is a set of measure zero in \mathfrak{R}^S .

We have shown existence of a set $\mathcal{G} \cup \mathcal{H}$ of Lebesgue measure zero in \mathfrak{R}^S with the property that if $\bar{\mathcal{U}}_i$ is singular, then $\mathcal{U}_i \in \mathcal{G} \cup \mathcal{H}$. By definition, then, W_i is an open set of full measure.

Part 2: *The matrix $J_{\mathcal{U}}[ii]$ defined in equation (26) has full row rank on W_i .*

To show this, 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 $[I + \delta \mathcal{P}_f^j] \cdot \bar{\mathcal{U}}_i$ has rank of S . Applying this fact again to the product of these two establishes that $J_{\mathcal{U}}[ii]$ has full rank of R_i . \square

Lemma 4 *For any $(f^0, \mathcal{U}^0) \in F^* \times W_1 \times W_2$ with W_1 and W_2 as in Lemma 3 and such that $M(f^0, \mathcal{U}^0) = 0$, there exist*

1. *a subset P_i of S for each i such that $|P_i| = |R_i|$ with corresponding notation $\mathcal{U}_P \equiv (\mathcal{U}_{1,P_1}, \mathcal{U}_{2,P_2})$ and $\mathcal{U}_Q \equiv (\mathcal{U}_{1,S \setminus P_1}, \mathcal{U}_{2,S \setminus P_2})$,*
2. *open neighborhoods G_F , G_Q and G_P of f^0 , \mathcal{U}_Q^0 and \mathcal{U}_P^0 , resp., with $G_Q \times G_P \subseteq W_1 \times W_2$, and*
3. *a locally smooth mapping $\Psi : G_F \times G_Q \rightarrow G_P \times G_Q$*

such that for each pair (f, \mathcal{U}_Q) , the vector \mathcal{U}_P given by $(\mathcal{U}_P, \mathcal{U}_Q) = \Psi(f, \mathcal{U}_Q)$ is the unique solution to $M(f, (\mathcal{U}_P, \mathcal{U}_Q)) = 0$.

Proof By Lemma 3, the submatrix $J_{\mathcal{U}}[ii]$ of $D_{\mathcal{U}}(M)$ 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 subset P_i of S such that $|P_i| = |R_i|$ with $\mathcal{U}_P = (\mathcal{U}_{1,P_1}, \mathcal{U}_{2,P_2})$ and $\mathcal{U}_Q = (\mathcal{U}_{1,S \setminus P_1}, \mathcal{U}_{2,S \setminus P_2})$ such that

$$\frac{\partial M}{\partial \mathcal{U}_{i,P_i}} \text{ has full rank } R_i \text{ at } (f^0, \mathcal{U}^0).$$

Letting G_F , G_Q and G_P denote the open neighborhoods of f^0 , \mathcal{U}_Q^0 and \mathcal{U}_P^0 , resp., with $G_Q \times G_P \subseteq W_1 \times W_2$, we apply the Implicit Function Theorem to obtain a locally smooth mapping

$\phi : G_F \times G_Q \rightarrow G_P$ satisfying $M(f, (\mathcal{U}_Q, \phi(f, \mathcal{U}_Q))) = 0$ for all $(f, \mathcal{U}_Q) \in G_F \times G_Q$. Moreover, for each $(f, \mathcal{U}_Q) \in G_F \times G_Q$, $\mathcal{U}_P = \phi(f, \mathcal{U}_Q)$ is the unique solution of $M(f, (\mathcal{U}_Q, \mathcal{U}_P)) = 0$ in G_P . The mapping $\Psi : G_F \times G_Q \rightarrow G_P \times G_Q$ defined by

$$\Psi(f, \mathcal{U}_Q) = (\phi(f, \mathcal{U}_Q), \mathcal{U}_Q) \quad (32)$$

therefore has the requisite properties. □

Now fix (f^0, \mathcal{U}^0) satisfying the properties of Lemma 4. Fix the neighborhoods G_F , G_P and G_Q and the mapping Ψ obtained in the Lemma. Call $(f; \mathcal{U}_Q) \in G_F \times G_Q$ an *f-regular point* of Ψ , if $\partial\Psi/\partial f$ has full rank $R_1 + R_2$ at $(f; \mathcal{U}_Q)$. Call $(f; \mathcal{U}_Q) \in G$ an *f-critical point* of Ψ otherwise. Call $(\mathcal{U}_P, \mathcal{U}_Q) \in G_P \times G_Q$ an *f-critical value* of Ψ , if $(\mathcal{U}_P, \mathcal{U}_Q) = \Psi(f; \mathcal{U}_Q)$ with $(f; \mathcal{U}_Q)$ an *f-critical point* of Ψ .

Lemma 5 *The set of f-critical values of Ψ has Lebesgue measure zero in $G_P \times G_Q$.*

Proof 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 (32). For any $\mathcal{U}_Q \in G_Q$, let $\phi_{\mathcal{U}_Q} : G_F \rightarrow G_P$ be defined by $\phi_{\mathcal{U}_Q}(f) = \phi(f, \mathcal{U}_Q)$ for $f \in G_F$. Then $\phi_{\mathcal{U}_Q}$ is a local smooth mapping with the properties

- (i) $M(f, (\mathcal{U}_Q, \phi_{\mathcal{U}_Q}(f))) = 0$ for all $f \in G_F$;
- (ii) For $f \in G_F$, $\mathcal{U}_P = \phi_{\mathcal{U}_Q}(f)$ is the only solution to the equation $M(f, (\mathcal{U}_Q, \mathcal{U}_P)) = 0$.

By Sard's theorem, the set of critical values of $\phi_{\mathcal{U}_Q}$ constitutes a set $G_P(\mathcal{U}_Q)$ of Lebesgue measure zero in G_P . In view of this finding, an application of Fubini's theorem suggests itself. However, this requires measurability of the set of critical values Z which we are going to demonstrate, following a lead in the proof of Sard's theorem in Milnor (1965), p. 18.

Let $G \equiv G_F \times G_Q$. The set $G \setminus Y$ is open relative to G and thus Y is closed relative to G . Let \mathbb{Q}^{2S} denote the set of points in \mathbb{R}^{2S} with rational coordinates and let \mathbb{Q}_{++} denote the set of strictly positive rational numbers. For a point $(x, r) \in \mathbb{Q}^{2S} \times \mathbb{Q}_{++}$, let $K_r(x)$ denote the closed ball in \mathbb{R}^{2S} with center x and radius r . For any $y = (f; \mathcal{U}_Q) \in G$, there exists $(x, r) \in \mathbb{Q}^{2S} \times \mathbb{Q}_{++}$ with $y \in K_r(x) \subset G$. Hence the closed set Y in G can be covered by a countable family of compact subsets $K_n, n \in \mathbb{N}$, of G . 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 $Z(\mathcal{U}_Q) = \{\mathcal{U}_P \in G_P : (\mathcal{U}_P, \mathcal{U}_Q) \in Z\}$ for each $\mathcal{U}_Q \in G_Q$. An element $\mathcal{U}_P \in G_P$ is a critical value for the mapping $\phi_{\mathcal{U}_Q}$, if there exists a critical point $f \in G_F$ of $\phi_{\mathcal{U}_Q}$ with $\mathcal{U}_P = \phi_{\mathcal{U}_Q}(f)$. But

$$\Psi(f; \mathcal{U}_Q) = (\phi(f; \mathcal{U}_Q), \mathcal{U}_Q) = (\phi_{\mathcal{U}_Q}(f), \mathcal{U}_Q)$$

Hence \mathcal{U}_P is a critical value of $\phi_{\mathcal{U}_Q}$ if and only if $(\mathcal{U}_P, \mathcal{U}_Q)$ is an f -critical value of Ψ . Therefore, $Z(\mathcal{U}_Q) = G_P(\mathcal{U}_Q)$ and, consequently, the section $Z(\mathcal{U}_Q)$ of Z has Lebesgue measure zero in G_P . Since this holds true for any choice of \mathcal{U}_Q and since Z is measurable, Fubini's Theorem implies that Z has Lebesgue measure zero in $G_P \times G_Q$. \square

Clearly, by Lemma 5 the f -regular values of Ψ have full measure in $G_P \times G_Q$.

Lemma 6 *Let $\mathcal{U} \in G_P \times G_Q$ such that \mathcal{U} is an f -regular value of Ψ . Then the elements of $G_F \cap \Gamma_{\mathcal{U}, C}$ are isolated points.*

Proof Recall the open sets G_F, G_P, G_Q , and the mapping $\Psi : G_F \times G_Q \rightarrow G_P \times G_Q$ with the measure zero set Z of f -critical values to which Lemma 5 applies. Let $(f, \mathcal{U}) = (f, \mathcal{U}_P, \mathcal{U}_Q) \in G_F \times G_P \times G_Q$ with $f \in \Gamma_{\mathcal{U}, C}$, $\mathcal{U} \notin Z$. Then we claim that f is an isolated point of $\Gamma_{\mathcal{U}, C}$.

Recall from Lemma 2 that $f \in \Gamma_{\mathcal{U}, C}$ is equivalent to $M(f, \mathcal{U}) = 0$. Hence $\mathcal{U}_P = \phi(f; \mathcal{U}_Q)$ and $(\mathcal{U}_P, \mathcal{U}_Q) = (\phi(f; \mathcal{U}_Q), \mathcal{U}_Q) = \Psi(f; \mathcal{U}_Q)$. It follows that $\mathcal{U} = (\mathcal{U}_P, \mathcal{U}_Q) \in \Psi(G)$. Then because of $\mathcal{U} \notin Z$, $(f; \mathcal{U}_Q)$ is an f -regular point of Ψ , i.e. $\partial\Psi/\partial f$ has full rank at $(f; \mathcal{U}_Q)$. Hence the derivative $D\phi_{\mathcal{U}_Q}$ has full rank at f . Therefore, by the Inverse Function Theorem, there exist open neighborhoods G_F^* of f in G_F and G_P^* of \mathcal{U}_P in G_P such that $\phi_{\mathcal{U}_Q}^*$, the restriction of $\phi_{\mathcal{U}_Q}$ to G_F^* , is a diffeomorphism from G_F^* onto G_P^* .

We conclude that f is the only point in $G_F^* \cap \Gamma_{\mathcal{U}, C}$. For suppose $g \in G_F^* \cap \Gamma_{\mathcal{U}, C}$. Then $g, f \in G_F^*$ and $M(g, (\mathcal{U}_Q, \mathcal{U}_P)) = M(f, (\mathcal{U}_Q, \mathcal{U}_P))$ where $\mathcal{U}_P \in G_P^*$. Therefore $\mathcal{U}_P = \phi_{\mathcal{U}_Q}^*(g) = \phi_{\mathcal{U}_Q}^*(f)$. Since $\phi_{\mathcal{U}_Q}^*$ is a diffeomorphism from G_F^* onto G_P^* , $g = f$ as asserted. \square

To summarize what we have shown so far: if $(f^0, \mathcal{U}^0) \in F^* \times W_1 \times W_2$ satisfies $M(f^0, \mathcal{U}^0) = 0$, then Lemmata 4-6 assert that there exist

1. an open neighborhood G_F of f^0 in F^* ;
2. an open neighborhood $G_P \times G_Q$ of \mathcal{U}^0 in $W_1 \times W_2$;
3. a subset Z of Lebesgue measure zero in $G_P \times G_Q$

such that for all $\mathcal{U} \in (G_P \times G_Q) \setminus Z$, the profiles $f \in G_F \cap \Gamma_{\mathcal{U}, C}$ are isolated points.

Without losing generality, we may assume that G_F is an open box or, more precisely, an open order interval in $\mathfrak{R}^{R_1+R_2}$ with rational valued end-points. Similarly, we assume $G_P \times G_Q$ to be an open box in \mathfrak{R}^{2S} , again with rational valued end-points. Consequently, application of Lemmata 4-6 to all qualifying pairs (f^0, \mathcal{U}^0) determines a sequence of triples $(G_F^n, G_P^n, G_Q^n, Z^n), n = 1, 2, \dots$, such that

- (a) for each pair $(f^0, \mathcal{U}^0) \in F^* \times W$, there exists some n with $(f^0, \mathcal{U}^0) \in G_F^n \times G_P^n \times G_Q^n$;
- (b) for each n , G_F^n is an open subset of F^* ;
- (c) for each n , $G_P^n \times G_Q^n$ is an open subset of $W_1 \times W_2$;
- (d) for each n , Z^n is a measure zero subset of $G_P^n \times G_Q^n$;
- (e) for each n , if $\mathcal{U} \in (G_P^n \times G_Q^n) \setminus Z^n$, then the elements of $G_F^n \cap \Gamma_{\mathcal{U}, C}$ are isolated points.

Now define for the carrier C ,

$$Z_C^\infty = (\mathfrak{R}^{2S} \setminus (W_1 \times W_2)) \cup \left(\bigcup_n Z^n \right)$$

Then Z_C^∞ has Lebesgue measure zero and for every $\mathcal{U} \in \mathfrak{R}^{2S} \setminus Z_C^\infty$, each element f of $\Gamma_{\mathcal{U}, C}$ is an isolated point, i.e. there exists an open neighborhood N_f of f such that $N_f \cap \Gamma_{\mathcal{U}, C} = \{f\}$.

Observe for each i , that the Markov strategies whose carrier is C_i or a subset thereof form a — typically lower-dimensional — non-empty, compact, and convex subset F_{i, C_i} of F_i . F_{i, C_i} is the product of sub-simplices of $\Delta(S_i)$ and contains $F_i(C_i)$. Denote $F_C = F_{1, C_1} \times F_{2, C_2}$. Recall from the definition of $\Gamma_{\mathcal{U}, C}$ that elements of $\Gamma_{\mathcal{U}, C}$ need not be Markov strategy profiles. Let $\hat{\Gamma}_{\mathcal{U}, C} = F_C \cap \Gamma_{\mathcal{U}, C}$. Each point $\hat{f} \in \hat{\Gamma}_{\mathcal{U}, C}$ is a Markov strategy profile which is isolated within F_C , i.e. there exists an open neighborhood $N_{\hat{f}}$ of \hat{f} relative to F_C such that $N_{\hat{f}} \cap \hat{\Gamma}_{\mathcal{U}, C} = \{\hat{f}\}$. This conclusion is valid as well in the cases where $R_1 = R_2 = \emptyset$. For then F_C is a singleton, i.e. of dimension 0, and the conclusion holds true trivially with $\hat{\Gamma}_{\mathcal{U}, C} = F_C$.

For the remainder of the proof let $\hat{\Gamma}_{\mathcal{U}} \equiv \bigcup_C \hat{\Gamma}_{\mathcal{U}, C}$ and suppose $\mathcal{U} \in \bigcap_C (\mathfrak{R}^{2S} \setminus Z_C^\infty)$ where the latter is a finite intersection.

Lemma 7 *Let $f \in \hat{\Gamma}_{\mathcal{U}}$. Then f is locally isolated in F .*

Proof Suppose $f \in \hat{\Gamma}_U$. Then $f \in \hat{\Gamma}_{U,C}$ for some carrier $C = (C_1, C_2)$. We have already established that f is an isolated point of $\hat{\Gamma}_{U,C}$ within F_C . But f could still be the limit of a sequence in $\hat{\Gamma}_U \setminus \hat{\Gamma}_{U,C}$. Suppose so. Then we can find a carrier $C' = (C'_1, C'_2)$ with $C' \neq C$ and can extract a sub-sequence $f^k, k \in \mathbb{N}$, such that $f^k \in \hat{\Gamma}_{U,C'}$ for all f^k . Then $f^k \rightarrow f$. It follows that f belongs to $\hat{\Gamma}_{U,C'}$, because of the appropriate choice of F^* in the case of carrier C' . Then, however, f is an isolated point of $\hat{\Gamma}_{U,C'}$ within $F_{C'}$. On the other hand, $f \in \hat{\Gamma}_{U,C}$ whereas $f^k \notin \hat{\Gamma}_{U,C}$ and thus $f^k \neq f$ for all k . This, however, combined with $f^k \rightarrow f$ contradicts the fact that f is an isolated point of $\hat{\Gamma}_{U,C'}$ within $F_{C'}$. Hence to the contrary, f is not the limit of a sequence in $\hat{\Gamma}_U \setminus \hat{\Gamma}_{U,C}$. It follows that f is an isolated point of $\hat{\Gamma}_U$ in F , as asserted. \square

For the given U , let E_U denote the set of MPE. E_U is a closed subset of F , hence compact. Further $E_U \subseteq \hat{\Gamma}_U$ and therefore by the previous Lemma the collection $N_{f^*}, f^* \in E_U$, is a covering of E_U by open, non-empty, pairwise disjoint sets. By compactness of E_U , there exists a finite sub-covering. But since each N_{f^*} contains no points in E_U other than f^* , this is only possible if E_U itself is finite. This concludes the proof. $\square\square$

References

- [1] Abreu, D. (1988), "On the Theory of Infinitely Repeated Games with Discounting," *Econometrica*, 56, 383-398.
- [2] 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.
- [3] Dutta, P. (1995), "A Folk Theorem for Stochastic Games," *Journal of Economic Theory*, 66, 1-32.
- [4] Dutta, P. (1996), "Efficient Markov Perfect Equilibria," mimeo, Columbia University, May.
- [5] Friedman, J.W. (1986), *Game Theory with Applications to Economics*. Oxford University Press: New York and Oxford.
- [6] Fudenberg, D. and E. Maskin (1986), "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information," *Econometrica*, 54, 533-56.
- [7] Fudenberg, D. and E. Maskin (1995), "Learning and Markov Behavior," mimeo, May.

- [8] Fudenberg, D., and J. Tirole (1991), *Game Theory*. The MIT Press: Cambridge, MA, and London.
- [9] Gilles, R.P., Haller, H.H., and P.H.M. Ruys (1997), "Semi-Core Equivalence", forthcoming in *Economic Theory*.
- [10] 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.
- [11] Harsanyi, J.C. (1971), "The Oddness of Equilibrium Points: A New Proof," *International Journal of Game Theory*, 2, 235-250.
- [12] Hildenbrand, W. (1974), *Core and Equilibria of a Large Economy*. Princeton University Press: Princeton.
- [13] Kalai, E., and W. Stanford (1988), "Finite Rationality and Interpersonal Complexity in Repeated Games," *Econometrica*, 56, 397-410.
- [14] Lagunoff, R. and A. Matsui (1997), "Asynchronous Choice in Repeated Coordination Games," forthcoming in *Econometrica*.
- [15] Mas-Colell, A. (1985), *The Theory of General Economic Equilibrium: A Differentiable Approach*. Cambridge University Press: Cambridge et al.
- [16] Maskin, E., and J. Tirole (1987), "A Theory of Dynamic Oligopoly, Part III: Cournot Competition", *European Economic Review* 31, 947-968.
- [17] Maskin, E. and J. Tirole (1988a), "A Theory of Dynamic Oligopoly, I: Overview and Quantity Competition with Large Fixed Costs," *Econometrica*, 56, 549-70.
- [18] Maskin, E. and J. Tirole (1988b), "A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand Curves and Fixed Costs," *Econometrica*, 56, 571- 600.
- [19] Milnor, J.W. (1965), *Topology From the Differentiable Point of View*. The University Press of Virginia: Charlottesville.
- [20] Wilson, R. (1971), "Computing Equilibria in N-Person Games," *SIAM Journal of Applied Mathematics*, 21, 80-87.