

Cowles Foundation Paper 873

Econometrica, Vol. 62, No. 4 (July, 1994), 745–781

STATIONARY MARKOV EQUILIBRIA

BY D. DUFFIE, J. GEANAKOPOLOS, A. MAS-COLELL, AND A. MCLENNAN¹

We establish conditions which (in various settings) guarantee the existence of equilibria described by ergodic Markov processes with a Borel state space S . Let $\mathcal{P}(S)$ denote the probability measures on S , and let $s \mapsto G(s) \subset \mathcal{P}(S)$ be a (possibly empty-valued) correspondence with closed graph characterizing intertemporal consistency, as prescribed by some particular model. A nonempty measurable set $J \subset S$ is *self-justified* if $G(s) \cap \mathcal{P}(J)$ is not empty for all $s \in J$. A *time-homogeneous Markov equilibrium* (THME) for G is a self-justified set J and a measurable selection $\Pi : J \rightarrow \mathcal{P}(J)$ from the restriction of G to J . The paper gives sufficient conditions for existence of compact self-justified sets, and applies the theorem: If G is convex-valued and has a compact self-justified set, then G has an THME with an ergodic measure. The applications are (i) stochastic overlapping generations equilibria, (ii) an extension of the Lucas (1978) asset market equilibrium model to the case of heterogeneous agents, and (iii) equilibria for discounted stochastic games with uncountable state spaces.

KEYWORDS: General equilibrium, stochastic games, ergodicity, overlapping generations, infinite-horizon economies.

INTRODUCTION

WE CONSIDER INFINITE HORIZON MODELS in which equilibrium is a matter of continual readjustment to the resolution of uncertainties in the form of “shocks” generated by an exogenous, time-homogeneous, Markov process. In this paper we search for equilibria that are dynamically simple in the following senses:

1. Equilibrium is a time-homogeneous stationary Markov process, where the current state is a sufficient statistic for the future evolution of the system. (The state space may include both exogenous and endogenous variables. The choice of state space will have consequences in the theory, and is a significant modeling choice in applications.)

2. The equilibrium process has an ergodic measure. (Roughly, an ergodic measure is a probability measure on the space of states such that the empirical distribution generated by the process converges to the ergodic measure almost surely, if started with the probability measure itself.)

We can give a variety of reasons (none original) to focus on such equilibria. They constitute the simplest sort of equilibria and are thus perhaps focal. There is also the suspicion that other equilibria require implausibly sophisticated coordination. To quote Guesnerie and Woodford (1992, Section 3), “an equilibrium which does not display minimal regularity through time—maybe stationarity—is unlikely to generate the coordination between agents that it assumes.” There is, finally, the fact that, whatever their merits, such restrictions are essential maintained hypotheses of asymptotic econometric methods.

¹ We are grateful for comments from Larry Blume, Hugo Hopenhayn, Sylvain Sorin, Jean-Francois Mertens, Richard Torres, Larry Jones, as well as the editor and referees. We are also grateful to the Mathematical Sciences Research Institute at Berkeley, where the first draft was completed in 1986.

Our main accomplishment is the elaboration of an abstract method of proving existence of the desired type of equilibrium. Briefly, we develop an approach whereby existence of equilibrium for a finite horizon version of a model implies that an “expectations correspondence” possesses certain properties that, in turn, imply the desired existence. The expectations correspondence is a temporary equilibrium device. It specifies, for each possible current state, the transitions (that is, distributions over states in the next period) that are consistent with feasibility *and* which, in conjunction with the current state, satisfy the short-run equilibrium conditions. Our results extend the work of Grandmont and Hildenbrand (1974), Hellwig (1980), and Blume (1982). These ideas are detailed in Section 1.

The remaining three sections illustrate our approach by applying it to problems in economics and game theory. As our first application of the central results, Section 2 of the paper demonstrates the existence of ergodic Markov equilibria for stochastic overlapping generations (OLG) economies with security markets for assets that are retraded period after period and pay dividends depending only on the current exogenous shock. Section 3 offers an extension of the asset market equilibrium model of Lucas (1978) to the case of heterogeneous agents. Given the influence of Lucas’ paper on financial economics, macroeconomics, monetary theory, and econometrics, it is significant that the existence of stationary Markov equilibria for this model does not rest on the assumption of a single type of agent. Section 4 of the paper demonstrates existence of equilibrium for noncooperative, nonzero sum stochastic games with uncountable state spaces. Our equilibria are “coordinated,” in the sense that they are given in terms of stationary Markov equilibria with respect to an extended state space that includes the players’ continuation values for the game.

1. THE CENTRAL THEORY

1.1. *Introduction*

We begin by describing the main elements of our approach (the state space and the “expectations correspondence”) in subsection 1.2, pointing out that the existence of a stationary Markov equilibrium process for such objects is a matter of finding a selection from the restriction (in both domain and range) of the expectations correspondence to a subset of the state space that has a property that we describe by the phrase “self-justification.” In subsection 1.3 we treat invariant and ergodic measures, supplying conditions under which there is a stationary Markov equilibrium process that admits an ergodic measure. Subsection 1.4 completes the formal description of our abstract methods by providing conditions under which existence of equilibrium for finite-horizon versions of the model can imply the existence of a self-justified set, so that earlier results can be applied.

In certain circumstances our existence results seem to require randomization over two or more possible equilibrium continuations from a given state, so that there is a dependence of the equilibrium process on uncertain events that are

outside the model as it is initially presented. Such dependence on “extrinsic uncertainty” is the defining feature of “sunspot equilibria” (Cass and Shell (1983)), a topic that has been studied extensively in recent years. In subsection 1.5 we provide a device for eliminating unnecessary dependence of transition probabilities on state variables. In subsection 1.6 we argue that the dependence on extrinsic randomness required by our methods is quite special, and mild in comparison with the possibilities considered by other authors.

1.2. *Time-Homogeneous Markov Equilibrium*

The first essential element of our approach is a nonempty Borel space² S of states. This is the space in which the equilibrium processes will live, and a state is to be thought of as a complete description of current variables (both exogenous and endogenous) that is, in addition, a sufficient statistic for the future evolution of the model. In many applications it will happen that $S = Y \times Z$ is the Cartesian product of a Borel space Y of exogenous shocks, which evolve according to an exogenously specified transition probability, and a Borel space Z of endogenous variables, but our general methods do not depend on such a decomposition.

Given S , the equilibrating forces in the model are described by an *expectations correspondence*, a map G from S to the subsets (including the empty set) of the space $\mathcal{P}(S)$ of probability measures on S , that has a closed³ graph. In the setting of a particular application, one would construct G so that a measure μ is in $G(s)$ if and only if: (i) μ is a distribution of tomorrow's state compatible with the constraints imposed by the current state (including consistency with the transition laws of exogenous shocks); and (ii) when the distribution of tomorrow's state is given by μ , no agent has an incentive to take an action different than that prescribed at s . A state must therefore contain enough information to allow all short-run equilibrium conditions to be verified by a comparison of s and $\mu \in \mathcal{P}(S)$.

A state space S and an expectations correspondence G are fixed for the remainder of Section 1.

The short-run equilibrium conditions embodied in G must be sufficient for optimality in the agents' dynamic programming problems. How this may occur is an underlying theme of the examples presented subsequently. For an overlapping generations economy in which all agents live for two periods, the optimization conditions can be expressed directly in terms of the comparison of today's state and the distribution of states for tomorrow, so there is no problem. In our treatment of the Lucas model of asset valuation, the appropriate transversality conditions follow from compactness conditions, so that the “Euler equation,”

² A Borel space is a measurable subset of a Polish (complete separable metric) space, endowed with the relative topology and the σ -algebra generated by relatively open sets.

³ Throughout, we give the space of probability measures on a Borel space the usual “weak-convergence” topology, under which $\mu_n \rightarrow \mu$ if and only if $\int f d\mu_n \rightarrow \int f d\mu$ for any bounded continuous real-valued f , and we define measurability with respect to the Borel σ -algebra.

which relates the marginal utility of consumption today and the expected marginal utility of consumption tomorrow, is not just necessary but also sufficient for optimality. In our treatment of discounted stochastic games, we resort to a technical device, expanding the state space by adding “continuation values.”

The least restrictive notion of equilibrium for S and G is a Markov process $\{s_1, s_2, \dots\}$ on some probability space such that, for all t , it is almost surely the case that the conditional distribution of s_{t+1} given $\{s_1, \dots, s_t\}$ is in $G(s_t)$. The idea of time-homogeneity is captured by the notion of a *transition*, a measurable map $\Pi: J \rightarrow \mathcal{P}(J)$ for some nonempty measurable $J \subset S$. For convenience, we write “ (J, Π) ” interchangeably with “ $\Pi: J \rightarrow \mathcal{P}(J)$,” and “ Π_s ” interchangeably with “ $\Pi(s)$.”

DEFINITION: *A time-homogeneous Markov equilibrium (THME) for G is a transition (J, Π) with $\Pi(s) \in G(s)$ for all $s \in J$.*

Given a THME (J, Π) , the usual theory of Markov processes (as in Revuz (1975)) shows how to construct an equilibrium in the sense just mentioned, that is, a time-homogeneous Markov process $\{s_1, s_2, \dots\}$ on some probability space such that, for all t , the conditional distribution of s_{t+1} given $\{s_1, \dots, s_t\}$ is almost surely $\Pi(s_t)$.

A necessary condition for the existence of THME is a *self-justified* set, a nonempty measurable set $J \subset S$ such that $G(s) \cap \mathcal{P}(J) \neq \emptyset$ for all $s \in J$. A self-justified set J merely allows for some way to draw the next period’s state from J for any current state in J . Intuitively (that is, disregarding issues of measurability), the existence of a self-justified set is also sufficient for existence. Indeed, if a self-justified J is closed (as it will typically be) then the point-to-set map G_J , defined by $G_J(s) = G(s) \cap \mathcal{P}(J)$, $s \in J$, always has a measurable selection $\Pi: J \rightarrow \mathcal{P}(J)$, by the Kuratowski-Ryll-Nardzewski Theorem (see Hildenbrand (1974, page 55)). We therefore have the simple result:

PROPOSITION 1.1: *If there is a closed self-justified set, then there is a time-homogeneous Markov equilibrium.*

It follows from Proposition 1.1 that the difficulty in applications will lie in establishing the existence of a self-justified set. This aspect of the theory is treated in Section 1.4.

1.3. Invariant Measures and Ergodic Markov Equilibria

There are respects in which the concept of a THME is too weak. In general, we are interested in a THME (J, Π) that is able to sustain some standard notions of stationarity. For example, an *invariant measure* for a THME (J, Π) , or indeed any transition $\Pi: J \rightarrow \mathcal{P}(J)$, is a measure $\mu \in \mathcal{P}(J)$ such that

$\mu = \mu\Pi$, where

$$\mu\Pi(A) = \int_J \Pi_s(A) d\mu(s).$$

That is, an invariant measure μ is defined by the property that, conditional on the information that the initial state is drawn with distribution μ , the distribution of the state at every future date is also μ . Not every THME (J, Π) has an invariant measure. For example,⁴ let $J = [0, 1]$, $\Pi_0 = \delta_1$, and $\Pi_s = \delta_{s/2}$, $s \in (0, 1]$.

A more restrictive notion of stationarity is given by an ergodic measure.

DEFINITION: If μ is invariant for a transition $\Pi: J \rightarrow \mathcal{P}(J)$, then a μ -invariant set is a measurable subset A of J satisfying $\Pi_s \in \mathcal{P}(A)$ for μ -a.e. $s \in A$. (For invariant μ , this condition implies that $\Pi_s \in \mathcal{P}(J \setminus A)$ for μ -a.e. $s \in J \setminus A$, and conversely.) An invariant measure μ is ergodic for (J, Π) if, for any μ -invariant set A , $\mu(A) = 0$ or $\mu(A) = 1$. If μ is ergodic for a THME (J, Π) , then (J, Π, μ) is an ergodic Markov equilibrium for G .

In stochastic settings ergodic, rather than invariant, measures are more natural analogues of the deterministic notion of steady state behavior. For example, let $S = \{0, 1\}$, $\Pi_0 = \delta_0$, and $\Pi_1 = \delta_1$. Then $\alpha\delta_0 + (1 - \alpha)\delta_1$, for any $0 < \alpha < 1$, is an invariant measure, but as soon as the process has started, this measure ceases to convey information about the future evolution of the system. This is not so for $\alpha = 0$ or $\alpha = 1$, which yield, precisely, the ergodic measures δ_0 and δ_1 .

Suppose μ is ergodic for (J, Π) . Let $\{s_t\}$ be a stationary Markov process in J induced by initial distribution μ and transition Π . A key property of ergodicity used in modern econometric theory (for example, Hansen (1982)) is the μ -almost-sure weak convergence of the empirical distribution of the state to the ergodic measure, that is:

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T h(s_t) = \int_J h d\mu,$$

for any $h \in L^1(J, \mu)$. This follows (if necessary, one should redefine the underlying probability space and make a particular choice of $\{s_t\}$ with the above properties) from the random ergodic theorem (Kifer (1986, Corollary 2.3)). By knowing that the average over time of a function h of the state process $\{s_t\}$ is the same as the average of h under the ergodic measure, one can obtain consistent estimators of parameters of the process.

Even if there exists an ergodic measure μ for a THME (J, Π) , it is not necessarily the case that, starting with an arbitrary distribution, either the actual or empirical distribution of the state converges to μ (a.s. with respect to the initial distribution). The significance of this question from a practical viewpoint

⁴ Here and subsequently δ_s denotes the Dirac measure at s , that is, the measure that assigns unit measure to the singleton $\{s\}$.

is clear, and an important thrust of statistical ergodic theory is to answer it affirmatively. The well known Döblin conditions and their generalizations (for example, Tweedie (1975)) are important here. The subject is deep, and we will not go into greater detail here.

For convenience, we say that μ is an *invariant measure* for a subset J of S if there exists a THME (J, Π) for which μ is an invariant measure. Our first theorem, characterizing the set of invariant measures of compact self-justified sets, is quite similar to Theorems 2.1 and 3.1 of Blume (1982). There the correspondence G is always nonempty valued and the state space S is assumed compact, but we place ourselves in essentially this setting in the statement of the theorem below by considering a compact self-justified set. (Later, we provide conditions for existence of a compact self-justified set that are easy to apply in practice.) Where we differ in approach is in the method of proof. Blume studies the properties of the correspondence $\bar{G}: \mathcal{P}(S) \rightarrow \mathcal{P}(S)$ defined by $\mu \in \bar{G}(\nu)$ if there exists a measurable selection Π from G such that $\mu = \nu\Pi$. A fixed point of \bar{G} is an invariant measure for S . We instead consider the “measurable selection problem” (that is, finding a THME) only after the invariant measure has been found. This seems to be simpler. We do not know whether the assumption that G is convex-valued is required in the following theorem. This assumption is discussed further in Section 1.5.

THEOREM 1.1: *If G is convex-valued and J is a compact self-justified set for G , then the set of invariant measures for J is a nonempty, compact, and convex subset of $\mathcal{P}(J)$.*

PROOF: Let G_J be the restriction of G to J in both domain and range: $G_J(s) = G(s) \cap \mathcal{P}(J)$, $s \in J$. Since $\text{Gr}(G)$ (the graph of G) is closed, $\text{Gr}(G_J)$ is compact. Let $m_1: \mathcal{P}[\text{Gr}(G_J)] \rightarrow \mathcal{P}(J)$ and $m_2: \mathcal{P}[\text{Gr}(G_J)] \rightarrow \mathcal{P}[\mathcal{P}(J)]$ be the restrictions to $\mathcal{P}[\text{Gr}(G_J)]$ of the functions that give the marginals of distributions on $J \times \mathcal{P}(J)$. For any $\eta \in \mathcal{P}[\mathcal{P}(J)]$, let $E\eta \in \mathcal{P}(J)$ denote its mean, which is uniquely defined, and weakly continuous in η , by the Riesz representation theorem for $C(J)$ (the continuous functions on J) and the definition of $E\eta$ as the unique measure with the property that, for any $f \in C(J)$, we have

$$\int_J f(s) dE\eta(s) = \int_{\mathcal{P}(J)} \int_J f(s) d\mu(s) d\eta(\mu).$$

Since J is self-justified, $m_1(\mathcal{P}[\text{Gr}(G_J)]) = \mathcal{P}(J)$, that is $m_1^{-1}(\mu) \neq \emptyset$ for all $\mu \in \mathcal{P}(J)$. (Clearly this is true if μ has finite support, and the set of such measures is dense in $\mathcal{P}(J)$ —see Hildenbrand (1974, page 49).) The functions m_1 , m_2 , and E are restrictions of continuous linear functions, so $E \circ m_2 \circ m_1^{-1}: \mathcal{P}(J) \rightarrow \mathcal{P}(J)$ is an upper semicontinuous convex-valued correspondence. Let M be its set of fixed points. Clearly M is compact and convex, and the Fan-Glicksberg fixed point theorem implies that M is nonempty.

We claim that each $\mu \in M$ is an invariant measure for some THME (J, Π) . Specifically, for $\mu \in M$ and $\nu \in \mathcal{P}[\text{Gr}(G_J)]$ with $Em_2(\nu) = m_1(\nu) = \mu$, there

exists a measurable $P: J \rightarrow \mathcal{P}[\mathcal{P}(J)]$ with $\mu \otimes P = \nu$, where $\mu \otimes P(A, B) \equiv \int_A P_s(B) d\mu(s)$. (The existence of P and of $\mu \otimes P$ are standard, and can be verified from Neveu (1965, Corollary (page 193) and Proposition III.2.1).) It need not be the case that $E \circ P: J \rightarrow \mathcal{P}(J)$ is a selection from G_J , but the Aumann measurable selection theorem (Hildenbrand (1974, p. 54)) implies the existence of a measurable selection from G_J , so a satisfactory Π can be constructed by replacing $E \circ P(s)$ on a μ -negligible set, once we show that $E \circ P(s) \in G_J(s)$ a.e. $[\mu]$.

Consider a bounded continuous $f: J \rightarrow \mathbb{R}$ and a constant $c \in \mathbb{R}$. Suppose that

$$\mu \left(\left\{ s \mid \max_{\rho \in G_J(s)} \int f d\rho \leq c < \int f d[E \circ P(s)] \right\} \right) > 0.$$

Let $A_{\mathcal{P}(J)} = \{\rho \in \mathcal{P}(J) \mid \int f d\rho > c\}$. Then $\mu(A_J) > 0$, where

$$A_J = \left\{ s \mid \max_{\rho \in G_J(s)} \int f d\rho \leq c \text{ and } P(s; A_{\mathcal{P}(J)}) > 0 \right\}.$$

Since $\text{Gr}(G_J) \cap (A_J \times A_{\mathcal{P}(J)}) = \emptyset$, we now have

$$0 = \nu(\emptyset) = \mu \otimes P(A_J \times A_{\mathcal{P}(J)}) = \int_{A_J} P(s; A_{\mathcal{P}(J)}) \mu(ds) > 0.$$

The upshot of this contradiction is that, for each $f \in C(J)$ (the space of bounded continuous real-valued functions on J),

$$\min_{\rho \in G_J(s)} \int f d\rho \leq \int f d[E \circ P(s)] \leq \max_{\rho \in G_J(s)} \int f d\rho, \quad \text{a.e. } [\mu].$$

Since the Banach space $C(J)$ is separable, there is a countable collection $\{f_j\} \subset C(J)$ such that $E \circ P(s) \in G_J(s)$ if this inequality holds for all f_j , so indeed $E \circ P(s) \in G_J(s)$ a.e. $[\mu]$. Q.E.D.

COROLLARY 1.1: *If G is convex-valued and there is a compact self-justified set, then there is an ergodic Markov equilibrium.*

PROOF: Let J be a compact self-justified set and μ be an extreme point of the set of invariant measures for J . Extreme points exist by the Krein-Milman theorem (see, for example, Royden (1968)). It remains to apply the definition of an extreme point to verify the properties of an ergodic measure for the transition function Π associated with μ . Suppose, with the goal of a contradiction, that A is a Π -invariant set and that $\mu(A) \in (0, 1)$. Then $\nu = \mu/[1 - \mu(A)]$ restricted to A^c (zero on A) and $\eta = \mu/\mu(A)$ restricted to A (zero on A^c) are both Π -invariant, and $\mu = \mu(A)\eta + [1 - \mu(A)]\nu$, which contradicts the fact that μ is extreme. Q.E.D.

Grandmont and Hildenbrand (1974) assume, in effect, that G has a continuous selection Π . In that case, the existence of an invariant measure follows

directly from Schauder's Theorem applied to the map $\mu \mapsto \mu\Pi$ from $\mathcal{P}(S)$ to $\mathcal{P}(S)$, for S compact. Our approach here, along with Blume (1982), is to obtain an invariant measure without requiring the existence of a continuous selection.

The potential multiplicity of equilibria should be a caution to the reader. We have not shown that *every* THME has an invariant measure. (Indeed, in the context of particular applications, there may generally exist equilibrium stochastic processes that do not have even the Markov property relative to a natural state space, let alone time-homogeneity, invariant measures, or ergodicity.) We have shown only that, given a compact self-justified set J , there is some particular selection from G_J that has an ergodic measure, which begs the issue of coordination among agents. Hellwig (1980) has gone on to show additional conditions in a temporary equilibrium context under which *any* selection from G_J has an invariant measure μ .

1.4. Existence of Self-Justified Sets

In some cases of interest, such as the temporary equilibrium model studied by Grandmont and Hildenbrand (1974), Hellwig (1980), and Blume (1982), the set S of states is compact and the expectations correspondence G is nonempty valued. In this case, S itself is self-justified and Theorem 1.1 can be applied directly. Going beyond the temporary equilibrium setting, our approach is to derive the possibility of ergodic behavior from the existence of finite-horizon equilibria. Specifically, given any time horizon $T \in \mathbb{N}$, a *T-horizon equilibrium* is an S -valued stochastic process $\{s_1, \dots, s_T\}$ on some probability space such that, for all $t < T$, the conditional distribution of s_{t+1} given $\{s_1, \dots, s_t\}$ is in $G(s_t)$ almost surely. (By "stochastic process," we mean simply that s_t is a random variable for all t .) A T -horizon equilibrium need not be Markov relative to the given state space S ; we will only use the fact that the conditional distribution of s_{t+1} given $\{s_1, \dots, s_t\}$ is in $G(s_t)$ almost surely. By the end of Part 1, we will have shown the following result, the easiest point of application of the Central Results.

PROPOSITION 1.2: *Let $G: S \rightarrow \mathcal{P}(S)$ be an expectations correspondence and K be a compact subset of S . If, for every $T \in \mathbb{N}$, there exists a T -horizon equilibrium $\{s_1, \dots, s_T\}$ for G such that $s_t \in K$ almost surely for all t , then G has an ergodic Markov equilibrium.*

In proving Proposition 1.2, we are proceeding in two basic steps. In the first, already accomplished, we have shown that the existence of a compact self-justified set is enough for the existence of an ergodic Markov equilibrium. This is the content of Theorem 1.1. We now proceed to the demonstration that the hypotheses of Proposition 1.2 imply the existence of a compact self-justified set.

Let K be a given subset of S . We begin by defining a set of states that contains the initial state of any " j -period equilibrium that stays in K ." Let

$C_{00} = K$, and for $j \geq 1$ define C_{0j} inductively by letting

$$C_{0j} = \{s \in K \mid \text{there exists } \nu \in G(s) \text{ with } \nu_*(C_{0,j-1}) = 1\}.$$

Here $\nu_*(C_{0,j-1})$ is the inner ν -measure of $C_{0,j-1}$, that is, the supremum of $\nu(D)$ over Borel measurable $D \subset C_{0,j-1}$. Induction immediately yields the following inclusion.

LEMMA 1.1: $C_{0j} \subset C_{0,j-1}$, $j \in \{1, 2, \dots\}$.

It will be a consequence of Theorem 1.2 that if K is compact and C_{0j} is not empty for all j , then $J = \bigcap_j C_{0j}$ is a self-justified set. We will formulate Theorem 1.2, however, so that it applies to cases in which K is noncompact.

Suppose that $i \geq 1$ and that $C_{i-1,j}$ has been defined for all j . For $j \in \{0, 1, 2, \dots\}$, let $C_{ij} = \{s \in C_{0j} \mid \text{there exists } s' \in C_{i-1,j+1} \text{ and } \nu \in G(s') \text{ with } \nu_*(C_{0j}) = 1 \text{ and } s \in \text{supp } \nu\}$. Roughly, C_{ij} contains all states that could occur in period i of an “ $(i + j + 1)$ -period equilibrium” in K running from period 0 to period $i + j$.

LEMMA 1.2: (a) $C_{ij} \subset C_{i-1,j}$ for all $i \geq 1$ and $j \geq 0$; (b) $C_{ij} \subset C_{i,j-1}$ for all $i \geq 0$ and $j \geq 1$.

PROOF: (a) When $i = 1$, this is an immediate consequence of the definition, and for larger i , it follows inductively from the inclusion $C_{i-1,j+1} \subset C_{i-2,j+1}$. (b) In view of the definition of C_{ij} , the inclusions $C_{0j} \subset C_{0,j-1}$ and $C_{i-1,j+1} \subset C_{i-1,j}$ imply that $C_{ij} \subset C_{i,j-1}$, so the result follows from Lemma 1.1 and induction on i . Q.E.D.

The next result confirms, at least to some extent, that we have defined the “correct” sets C_{ij} . In addition, it simplifies the verification of the hypotheses of Theorem 1.2 below.

LEMMA 1.3: If $i \geq 1$ and $j \geq 0$, then $C_{ij} \neq \emptyset$ if and only if $C_{i-1,j+1} \neq \emptyset$.

PROOF: That $C_{ij} \neq \emptyset$ implies $C_{i-1,j+1} \neq \emptyset$ is implicit in the definition of C_{ij} . Consequently, it suffices to show that $C_{0j} \neq \emptyset$ implies $C_{i,j-i} \neq \emptyset$ for $i \in \{1, \dots, j\}$.

Fix $s_0 \in C_{0j}$, choose $\nu_1 \in G(s_0)$ with $\nu_{1*}(C_{0,j-1}) = 1$, and choose $s_1 \in C_{0,j-1} \cap \text{supp } \nu_1$. In general, if $s_i \in C_{0,j-i}$ has already been chosen, then we choose $\nu_{i+1} \in G(s_i)$ with $\nu_{i+1*}(C_{0,j-i-1}) = 1$ and $s_{i+1} \in C_{0,j-i-1} \cap \text{supp } \nu_{i+1}$. The definition of C_{ij} now implies, by induction, that $s_i \in C_{i,j-i}$ for all i . Q.E.D.

Our second main result is as follows. Note that it does not require the assumption that G is convex-valued.

THEOREM 1.2 (Existence of Self-Justified Sets): *Suppose that G is an expectations correspondence for which all C_{ij} are nonempty. (By Lemma 1.3, it suffices that $C_{0j} \neq \emptyset$ for all j .) Fixing an i^* , if C_{i^*j} has compact closure for some j , then $J = \bigcap_j \bar{C}_{i^*j}$ is self-justified.*

PROOF: Since J is the intersection of a nested sequence of nonempty compact sets, J is nonempty and compact. Fix $s \in J$ and choose a sequence $\{s_j\}$ with $s_j \in C_{i^*j}$ and $s_j \rightarrow s$.

For each j , choose $v_j \in G(s_j)$ with $v_j(C_{0,j-1}) = 1$; this is possible since $s_j \in C_{0j}$ by Lemma 1.2. Then

$$C_{0,j-1} \cap \text{supp } v_j \subset C_{i^*+1,j-1}$$

by virtue of the definition of C_{i^*+2j-1} , so

$$v_j \in \mathcal{P}(\bar{C}_{i^*+1,j-1}) \subset \mathcal{P}(\bar{C}_{i^*,j-1}).$$

Now observe that $\{\mathcal{P}(\bar{C}_{i^*j})\}$ is a descending sequence of compact sets with intersection $\mathcal{P}(J)$, so there is a limit point $\nu \in \mathcal{P}(J)$ of the sequence $\{v_j\}$. The fact that $\nu \in G(s)$ now follows from the assumption that the graph of G is closed. *Q.E.D.*

Proposition 1.2, which, to repeat, is the easiest point of application of our central results, can now be viewed as corollary of Theorems 1.1 and 1.2. First, the hypotheses of Proposition 1.2 imply (by induction in T) that C_{0T} is not empty for all T , taking C_{00} to be the compact set K referred to in Proposition 1.2. Theorem 1.2 then implies the existence of a compact self-justified set J , and the conclusion of Proposition 1.2 then follows from Theorem 1.1. The combination of Theorems 1.1 and 1.2 thus completes our central theory.

Roughly speaking, what is going on is the following. We wish to obtain the existence of a stochastic process whose transition is characterized as an ergodic Markov equilibrium. This is hard to do by a direct fixed point argument on a big space (say, of endogenous stochastic processes) because, while we may get existence, there is no guarantee about ergodicity, or even time-homogeneous Markov transitions: The fixed point in the big space neglects to use the inner recursivities of the problem. We therefore proceed by separating the construction into two parts. The first (which is model-specific and constitutes the input for Theorem 1.2) carries out a fixed point proof in a space of arbitrarily large, but finite, trajectories. Stationarity considerations are neglected at this stage. Only one thing is finally retained of the construction of this first part, namely the existence of a compact self-justified set (Theorem 1.2). The second part proceeds then, so to speak, to stationarize the equilibrium trajectories obtained

in the first part. This is done by plugging the self-justified set into Theorem 1.1 and getting a THME with an ergodic pressure.

1.5. *Elimination of Unnecessary Dependence on the Past*

In many applications the expectations correspondence may depend on a proper subset of the information presented by elements of S . An example is the overlapping generations model considered in Section 2: if s is the state today, then $G(s)$ does not depend on the components of s representing the initial portfolio of the current old. We now show that our results may be applied in a way that produces a THME, with all the properties derived above, in which the transition does not depend on “irrelevant” information. That the distinction may have consequences is shown by Spear, Srivastava, and Woodford (1990), who consider a framework in which equilibria restricted to depend only on “relevant” information are deterministic, in the sense of local uniqueness, but unrestricted equilibria are not.

Assume that \tilde{S} is a Polish space, $Q: S \rightarrow \tilde{S}$ is a continuous surjection (think of Q as a projection), and $\tilde{G}: \tilde{S} \rightarrow \mathcal{P}(S)$ is an upper semicontinuous set-valued mapping such that $G = \tilde{G} \circ Q$. We apply our results to \hat{S} and \hat{G} , where $\hat{S} = S \cup \tilde{S}$ (with the disjoint union topology) and $\hat{G}: \hat{S} \rightarrow \mathcal{P}(\hat{S})$ is given by letting $\hat{G}(\hat{s}) = Q(\hat{s})$ if $\hat{s} \in S$ and letting $\hat{G}(\hat{s}) = \tilde{G}(\hat{s})$ if $\hat{s} \in \tilde{S}$. The key facts (which are easily proved by straightforward arguments) are as follows:

(•) $J \subset S$ is self-justified for G if and only if $\hat{J} = J \cup Q(J)$ is self-justified for \hat{G} , and J is compact if and only if \hat{J} is.

(•) If $(\hat{J}, \hat{\Pi})$ is a THME for \hat{G} , then (J, Π) is a THME for G if we set $J = \hat{J} \cap S$ and $\Pi = \hat{\Pi} \circ Q|_J$.

(•) With $(\hat{J}, \hat{\Pi})$ and (J, Π) as in the last remark, if $\hat{\mu}$ is an ergodic measure for $(\hat{J}, \hat{\Pi})$, then $\mu \in \mathcal{P}(J)$ given by $\mu(A) = 2\hat{\mu}(A)$ is an ergodic measure for (J, Π) .

Given a self-justified set for G , the first remark above implies that our results above may be applied to $\hat{J} = J \cup Q(J)$; the second remark provides a procedure for “translating” the THME obtained for \hat{G} into an equilibrium for G in which the transitions do not depend on the irrelevant information; and the third remark shows that the existence of an ergodic measure is preserved by this procedure.

1.6. *Sunspots*

Our approach to the existence of ergodic measures does not use the full power of the additional structure on G that may often be present in applications. Suppose, for example, that $S = [0, 1]$ and for each $s \in S$, $G(s) = \delta_{g(s)}$ where $g: [0, 1] \rightarrow [0, 1]$ is a continuous deterministic transition. This model admits a deterministic ergodic measure (that is, a steady state), namely $\delta_{\bar{s}}$, where \bar{s} is any fixed point of g . Yet the ergodic measures obtained from

Corollary 1.1 need not be deterministic; they may capture cycles or other recurrent behavior of g . This is because, even if the underlying deterministic model belongs to a class having a steady-state, this information is never used in the construction of our central theory. In this sense, our construction is not “best,” which may be a partial explanation of why the interpretation of our equilibrium requires some notion of sunspots.

The phrase “sunspot equilibrium” (Cass and Shell (1983)) describes a situation in which equilibrated variables respond to random events that have no effect on the fundamentals of the economy—tastes, endowments, and technology. Implicitly, this issue has already arisen in connection with the convex-valuedness of G . As we have seen, given a compact self-justified set, one can construct an THME even if G is not convex-valued, but in order to guarantee the existence of an invariant measure, we need to convexify; that is, it may be necessary to allow THME’s that randomize over “pure” continuations if the given expectations correspondence is not already convex-valued.

The above sunspots, however, are trivial in the sense of Cass and Shell (1983). For any $s \in S$, each possible “pure” continuation μ in $G(s)$ represents a set of expectations about the future for which the choice variables in s are optimal, so that (given von Neumann-Morgenstern preferences) a randomization over such continuations is in turn a continuation for which the currently equilibrated variables are optimal. Among other things, one could imagine that before agents make their choices in the current period, the randomization over continuations has already taken place, and its results are known to all.

The issue of sunspots has a richer significance when one can meaningfully distinguish between “intrinsic” and “extrinsic” uncertainty. Suppose, for example, that $S = Y \times Z$ is the Cartesian product of a Borel space Y of exogenous shocks and a Borel space Z of endogenous variables. Suppose the variables in Y evolve according to an exogenously given stationary Markov process $\{y_t\}$ with a unique invariant measure. It would be natural to assume that the uncertainty in Y is purely intrinsic and payoff relevant: Different elements of Y correspond to different fundamentals.

Suppose that $\Pi: J \rightarrow \mathcal{P}(J)$ is a THME. Skorokhod’s theorem (Billingsley (1986, p. 343)) allows us to represent the $(Y \times Z)$ -valued state process $\{s_t = (y_t, z_t)\}$ in the form $z_{t+1} = f(y_{t+1}, \alpha_{t+1}, s_t)$, where $\alpha_{t+1} \in [0, 1]$ is uniformly distributed and *i.i.d.* We can think of α_t as the sunspot at time t . It would be of interest to have equilibria in which z_{t+1} does not depend on the sunspot α_{t+1} . This corresponds to our next definition.

DEFINITION: Let $\mathcal{P}_F(Y \times Z)$ denote the set of μ in $\mathcal{P}(Y \times Z)$ for which there is some $h: Y \rightarrow Z$ with measurable graph such that $\mu(\text{Gr}(h)) = 1$. A THME $\Pi: J \rightarrow \mathcal{P}(J)$ is *spotless* if $\Pi_s \in \mathcal{P}_F(Y \times Z)$ for all s in J .

Given a spotless THME, the distribution of the next period’s state supports a unique endogenous variable for each exogenous variable, and therefore sunspots

are absent. In general, we cannot prove the existence of a spotless THME, although additional conditions allow us to obtain the following substitute.

DEFINITION: A THME $\Pi: J \rightarrow \mathcal{P}(J)$ is conditionally spotless if, for all $s \in J$, there is some $M \subset \mathcal{P}_F(Y \times Z) \cap G(s)$ and $\pi \in \mathcal{P}(M)$ such that $\Pi_s = \int \nu d\pi(\nu)$.

With a conditionally spotless THME, sunspots are used only to randomize over spotless transitions to the next period's state. In particular, any optimality properties of agents' decisions that are used to define $G(s)$ are still (almost surely) satisfied even if the agent were to condition on the outcome of the sunspot governing the choice of the next period's state. Put another way, the state process is such that its dependence on the underlying sunspot process could (but need not) have the form $z_{t+1} = f(y_{t+1}, \alpha_t, s_t)$ in the sense that the dependence of z_{t+1} on sunspots is restricted to those sunspots that have already become publicly observable at the time the choice variables of s_t are selected.

We should emphasize that the existence of a conditionally spotless equilibrium is a nontrivial condition that will not be satisfied by arbitrary expectations correspondences. For example, if $Y = \{0\}$, $Z = \{0, 1\}$, and for each s , $G(s)$ contains a single stochastically nontrivial transition, then for each $s \in S$, $\mathcal{P}_F(Y \times Z) \cap G(s) = \emptyset$, so a conditionally spotless equilibrium is impossible.

We can, however, demonstrate conditionally spotless THME with an ergodic measure under certain conditions.

PROPOSITION 1.3: Suppose Y is finite, and for $s \in S$ let

$$g(s) = \{\nu \in G(s) \mid \nu \in \mathcal{P}_F(Y \times Z)\}.$$

Then g is an expectations correspondence. If g has a compact self-justified set J , then g has a spotless THME, and there is an ergodic Markov equilibrium for the expectations correspondence $\hat{G}: J \rightarrow \mathcal{P}(J)$ given by letting $\hat{G}(s)$ be the closure of the convex hull of $g(s) \cap \mathcal{P}(J)$. If, moreover, G is convex-valued, then an ergodic Markov equilibrium for \hat{G} is a conditionally spotless ergodic Markov equilibrium for G .

PROOF: Given that Y is finite, the graph of g is a closed subset of the graph of G , hence closed since the graph of G is closed. Thus g is an expectations correspondence. Given a compact self-justified set J for g , a THME for g is given by measurable selection from its restriction g_J to J , as discussed earlier.

We claim that \hat{G} has a closed graph. To show this let s_k and μ_k be sequences with $\mu_k \in G(s_k)$ for all k , $s_k \rightarrow s$, and $\mu_k \rightarrow \mu$. For each k there is $\nu_k \in \mathcal{P}(g(s_k)) \subset \mathcal{P}(\mathcal{P}(J))$ such that μ_k is the "mean" of ν_k . Since J is compact, so is $\mathcal{P}(\mathcal{P}(J))$, so ν_k has a convergent subsequence, say with limit ν . Elementary properties of the weak* topology imply that $\text{supp } \nu \subset g(s)$ and that μ is the mean of ν , so $\mu \in G(s)$.

If J is self-justified for g , then it is self-justified for \hat{G} , of course, so that \hat{G} has an ergodic Markov equilibrium by Corollary 1.1. The final assertion follows simply from the fact that, when G is convex-valued, $\hat{G}(s) \subset G(s)$ for all s .
Q.E.D.

The starting point of the analysis of a particular application will often be a result (for example, Radner (1972)) guaranteeing the existence, for any initial condition and finite horizon, of an equilibrium in which all temporary equilibria are drawn from g . We saw in the last section how to pass from this sort of result to the existence of a self-justified set for g . This means, quite generally, that if Y is finite and g has a self-justified set, then there exists a spotless THME.

In order to obtain an ergodic measure, however, it is necessary to convexify g , so we end up with a conditionally spotless equilibrium. The convexification of g seems quite natural, but on the other hand we know of no example satisfying our assumptions (with a finite Y) in which there is a self-justified set for g that does not admit a spotless THME with an ergodic measure. We wish to stress that the question of whether such a situation can exist is an open problem of considerable interest: a positive answer would suggest that our methods are capable of extracting all available general results, so that little remains unsaid, while a negative answer would, in all likelihood, involve novel and interesting methods of analysis.

Insofar as there is typically some latitude in the choice of the state space, the Markov property embodied in the notion of a THME has considerable flexibility in practice. Among other things, one can always replace S with $S^- \times S$, where S^- is a copy of S , to be thought of as a record of the last period's state. Since the expectations correspondence already embodies the equilibrium conditions, allowing such dependence on the past is similar, in a particular sense, to allowing sunspots: prior states are no longer "fundamental." The strongest (in the sense of minimizing dependence on the past) existence results are obtained from our methods by specifying a minimal state space. However, this does not mean that it cannot happen that, by allowing more than minimal dependence on the past, one can obtain equilibria with other desirable properties. Among other things, we are unable to rule out the possibility of an example with no spotless ergodic equilibria that has spotless ergodic equilibria when the state space is expanded to include lagged variables; whether this can happen is another interesting open problem. In passing, it should be stressed that in some applications (such as repeated games which, in the terminology of Section 4, are stochastic games with trivial state space) the interesting equilibria may well be those with more than minimal dependence on the past.

In many applications the coordination problems posed by our equilibrium concepts are severe, and they certainly become more problematic when the randomization discussed above is introduced. Concretely, it requires that the agents agree that an irrelevant random variable will determine which of a number of equally valid continuations will be followed. As Chiappori and Guesnerie (1992) make clear by example, once one allows arbitrary sorts of

sunspot equilibria, the multiplicity of equilibria and associated coordination issues become even more severe.

2. MARKOV OVERLAPPING GENERATIONS EQUILIBRIA

2.1. Introduction

Since Samuelson (1958) described the consumption-loan model, overlapping generations economies have been at the center of microeconomic theories about macroeconomic phenomena. (See, for example, Sargent (1987) and Geanakoplos (1987) for general accounts.) This is true both of the deterministic and, to a lesser extent, also of the stochastic version of the theory (as in, for example, Peled (1985)).

In this part we concern ourselves with the existence of Markovian stationary rational expectations equilibrium for stochastic overlapping generations economies with many households and commodities. This question has already been studied by Spear and Srivastava (1986) as well as Spear, Srivastava, and Woodford (1990). We establish existence of ergodic Markov equilibria in this setting as a straightforward application of our central results. An extension to a model with production has been given by Wang (1990).

This section is in some respects the most elementary in the paper because the optimality of an agent's behavior in a THME is an immediate consequence of the short run equilibrium conditions embodied in the definition of the expectations correspondence. In the applications studied in subsequent sections the relationship between equilibrium and optimization will be less straightforward.

2.2. The Primitives

The "fundamentals" of the economy have the following description.

There is a random process $\{y_t\}$ of exogenous shocks, a time-homogeneous Markov chain with finite state space Y and transition $P: Y \rightarrow \mathcal{P}(Y)$.

The number of perishable consumption commodities (in each period) is l , and for agents of type i the set of feasible consumption in a single period is $X_i = \mathbb{R}_+^l$. The number of agent types (in each generation) is m , and the number of (infinitely lived) assets is n .

Each agent lives for 2 periods. An agent of type i is defined by (e_i^1, e_i^2, u_i) , where:

(i) $e_i^1: Y \rightarrow \mathbb{R}_{++}^l$ is the endowment function for the first period of the agent's life: that is, when young, an agent of type i receives the bundle $e_i^1(y)$ if the current shock is y .

(ii) $e_i^2: Y \rightarrow \mathbb{R}_{++}^l$ gives the endowments of agents of type i , as a function of the shock y , when old.

(iii) $u_i: X_i \times X_i \times Y \times Y \rightarrow \mathbb{R}$ is a von Neuman-Morgenstern utility function for agents of type i , with $u_i(x_i^-, x_i, y^-, y)$ representing the utility received over the total lifetime of the agent if x_i^- is consumed when young and x_i is consumed

when old, while the shock is y^- when young and y when old. For each $(y^-, y) \in Y^2$, the function $u_i(\cdot, \cdot, y^-, y)$ is assumed to be increasing, continuous, and concave.

The infinitely lived securities are “real” in the sense that they generate commodity (as opposed to nominal) dividends. The dividend structure is specified by $d: Y \rightarrow (\mathbb{R}_+^l)^n$, with $d_j(y) \in \mathbb{R}_+^l$ denoting the dividend bundle paid by security j whenever the current shock is y . In each period the old realize the dividends generated by their portfolios given by the current shock, after which the securities are sold ex dividend to the young. By way of an innocuous normalization, the total endowment of each security is taken to be 1.

The economy described above is denoted (P, e, u, d) .

2.3. Equilibrium Definition

We now define spaces in which the variables describing economic activity live. The price simplex is

$$\Delta^{l+n} = \left\{ (p, q) \in \mathbb{R}_+^n \times \mathbb{R}_+^l \mid \sum_{j=1}^n p_j + \sum_{h=1}^l q_h = 1 \right\},$$

with p representing security prices and q representing commodity prices.

The space of possible vectors of portfolios held by a generation of young agents at the end of their first period of life is

$$\Lambda = \left\{ \theta = (\theta_1, \dots, \theta_m) \in (\mathbb{R}_+^n)^m \mid \sum_{i=1}^m \theta_i = (1, \dots, 1) \right\}.$$

Let $X = X_1 \times \dots \times X_m$ be the space of consumption profiles for a single generation in a single period. Let X^1 and X^2 be copies of X , interpreted as the spaces containing the consumption profiles of the young and old respectively. Let Λ^1 and Λ^2 be copies of Λ , interpreted as the spaces containing, respectively, the final portfolio of the current young and the initial portfolio of the current old. The “endogenous state space” is

$$Z = \Lambda^1 \times \Lambda^2 \times X^1 \times X^2 \times \Delta^{l+n},$$

consisting of portfolios, feasible consumptions, and prices. The set of states, that is, possible configurations of exogenous and endogenous states, is

$$S = \left\{ [y, (\theta^1, \theta^2, x^1, x^2, p, q)] \in Y \times Z \mid \sum_{i=1}^m [x_i^1 + x_i^2 - e_i^1(y) - e_i^2(y)] = \sum_{j=1}^n d_j(y) \right\}.$$

DEFINITION 2.1: A state transition for (P, e, u, d) is a pair (J, Π) , where J is a nonempty measurable subset of S and $\Pi: J \rightarrow \mathcal{P}(J)$ is a transition function. A state transition (J, Π) is consistent if, for each $s = [y, (\theta^1, \theta^2, x^1, x^2, p, q)] \in J$:

- (a) the marginal of $\Pi(s)$ on Y is $P(y)$;
- (b) the marginal of $\Pi(s)$ on Λ^2 is the dirac measure at θ^1 .

Condition (a) means that Π is consistent with the exogenous shock transition function P . Condition (b) says that the initial portfolio of the old tomorrow must equal the final portfolio of the young today.

From the point of view of a currently young agent, the decision-relevant economic data consists of the current state of the economy and the conditional distribution over states in the subsequent period. We have the following terminology for the agent's planning problem.

DEFINITION 2.2: A plan for an agent of type i is a triple $(\hat{x}_i^1, \hat{\alpha}_i, \hat{\xi}_i)$, where $\hat{x}_i^1 \in X_i$, $\hat{\alpha}_i \in \mathbb{R}_+^n$, and $\hat{\xi}_i: S \rightarrow X_i$ is a measurable function. A plan $(\hat{x}_i^1, \hat{\alpha}_i, \hat{\xi}_i)$ is budget-feasible, relative to a given pair $(s, \mu) \in S \times \mathcal{P}(S)$, if

$$(2.1) \quad p \cdot \hat{\alpha}_i = q \cdot [e_i^1(y) - \hat{x}_i^1],$$

where p, q , and y are the respective components of s , and

$$(2.2) \quad p(s') \cdot \hat{\alpha}_i + q(s') \cdot \left[e_i^2(y(s')) + \sum_{j=1}^n \hat{\alpha}_{ij} d_j(y(s')) \right] = q(s') \cdot \hat{\xi}_i(s'),$$

for all μ -a.e. $s' \in S$. A plan $(\hat{x}_i^1, \hat{\alpha}_i, \hat{\xi}_i)$ is optimal relative to (s, μ) if it is budget-feasible and, for any other budget-feasible plan $(\tilde{x}_i^1, \tilde{\theta}_i, \tilde{\xi}_i)$,

$$(2.3) \quad E(u_i[\hat{x}_i^1, \hat{\xi}_i(s'), y(s), y(s')]) \geq E(u_i[\tilde{x}_i^1, \tilde{\xi}_i(s'), y(s), y(s')]),$$

where the expectations are taken with s' having distribution μ .

An equilibrium is a consistent state transition in which the demands of the various agents specified by the corresponding components of the states are, in fact, optimal, in the following sense.

DEFINITION 2.3: For each state $s = [y, (\theta^1, \theta^2, x^1, x^2, p, q)]$, the derived plan of an agent of type i is $(x_i^1, \theta_i^1, \text{proj}_{X_i^2})$, where $\text{proj}_{X_i^2}: S \rightarrow X_i^2$ is the natural projection.

DEFINITION 2.4: An equilibrium transition for (P, e, u, d) is a consistent state transition (J, Π) such that, for each i and each $s \in J$, the derived plan of an agent of type i is optimal for an agent of type i given $(s, \Pi(s))$.

Note that the first budget constraint (2.1) is embedded in the requirement that the derived plans are optimal (hence budget-feasible) at each $s \in J$. The second budget constraint (2.2) need not hold at all $s \in J$, but it must hold at all

points in the support of any $\Pi(s)$. Our method of expressing optimization for an agent of type i does not call for optimality of consumption relative to the situation of an agent of type i when old. However, a plan that is optimal in our sense (that is, for the agent when young) remains optimal (with probability one) after the passage of time and revelation of information.⁵

Time-consistency of preferences is implicitly assumed in the sense that each agent has only one utility function, rather than one utility function when young and another when old, in which case the relevant notion of optimality would be to have the young acting as Stackelberg leaders for themselves when old. There is no apparent reason that our results could not be generalized in this direction.

We will be interested in demonstrating the existence of an equilibrium transition with several properties: ergodicity, conditional spotlessness, and compatibility with arbitrary initial conditions. In general it is desirable that an equilibrium transition be compatible with arbitrary initial conditions since, in applications, it may be helpful to have a continuation of an equilibrium from any reasonable starting point. In the current model, the historically determined information consists of the preferences and portfolios of the old.

DEFINITION 2.5: *An equilibrium transition (J, Π) is complete if, for any continuous, concave, strictly increasing utility functions $v_1, \dots, v_m: Y \times \mathbb{R}_+^l \rightarrow \mathbb{R}$, and $\hat{y} \in Y$, and any $\hat{\alpha}^2 \in A^2$, there is a state $s = [y, (\theta^1, \theta^2, x^1, x^2, p, q)] \in J$ such that $y = \hat{y}$, $\theta^2 = \hat{\alpha}^2$, and for each agent of type i , x_i^2 is a solution to the problem of maximizing $v_i(\hat{y}, \cdot)$ subject to the budget constraint*

$$q \cdot x_i^2 = p \cdot \hat{\alpha}_i^2 + q \cdot \left[e_i^2(\hat{y}) + \sum_{j=1}^n \hat{\alpha}_{i,j}^2 d_j(\hat{y}) \right].$$

We are ready to state the main result of this section.

THEOREM 2.1: *For any economy (P, e, u, d) , there exists a complete equilibrium transition with an ergodic measure.*

As remarked at the end of this section, our method of proof actually implies the existence of conditionally spotless equilibrium, in the sense described in subsection 1.6.

2.4. Proof of Theorem 2.1

The proof of Theorem 2.1 is a rather simple application of our central results. We begin by defining an expectations correspondence. Let $g: S \rightarrow \mathcal{P}(S)$ be the

⁵ It has been traditional, at least since Debreu (1953), not to formally state optimality conditions beyond those governing the initially chosen plan.

set-valued function defined by $(s, \mu) \in \text{Gr}(g)$ if and only if (with $s = [y, (\theta^1, \theta^2, x^1, x^2, p, q)]$):

(i) the support of μ is the graph of a function $h: Y \rightarrow \Lambda^1 \times \Lambda^2 \times X^1 \times X^2 \times \Delta^{l+n}$;

(ii) the marginal of μ on Y is $P(y)$;

(iii) the marginal of μ on Λ^2 is δ_{θ^2} ;

(iv) for each agent of type i , the associated derived plan of agent i for state s is optimal given (s, μ) .

We induce a T -period equilibria for all finite horizons $T \in \mathbb{N}$ as follows. Let the preferences of the “old at time zero” and the “young at time T ” be arbitrary. For each T these preferences induce a T -period finite-state Radner (1972) style event-tree economy. Then, by an argument similar to Radner’s, we establish the existence of an S^T -valued random variable that is a T -horizon equilibrium for G , in the sense of our central results. (Radner’s theorem fails to apply only by virtue of minor technical differences between our framework and his; a proof applicable to our framework is available upon request.) Since the amount of each good available in each state and date is bounded uniformly, both above and away from zero, there is a compact set K almost surely containing all states reached in any such T -horizon equilibrium. Proposition 1.2 now implies the existence of a compact self-justified set J for g , and the desired conclusion follows from Proposition 1.3.

2.5. Informal Discussion

As our results are formulated above, the distribution of states in the next period is allowed to depend on all the variables in the current state, even though the only variables that are obviously relevant to equilibration in the next period are the current shock, the consumptions of the current young, and the final portfolios of the current young. As shown in Section 1.5, it is possible to prove the existence of equilibria with the properties indicated in Theorem 2.1 in which, whenever the variables just enumerated are the same, the distribution of states in the next period is the same. Since Y is finite, the equilibrium demonstrated is also conditionally spotless in the sense described in Section 1.6.

Our stochastic setting allows a richer set of questions than would be possible for a simpler model. Consider an example in which there is an asset that pays no dividend in every state. Should such an asset have a nonzero price in equilibrium, it could be described as “money,” in the sense of a “fundamentally” valueless store of value, or as a “bubble.” Samuelson (1958) showed that this is a possibility for deterministic overlapping generations economies, but Scheinkman (1980) later established that such bubbles cannot coexist with assets paying positive dividends: for both assets to be held, the capital gain on the bubble must match the return on the dividend-paying asset, leading eventually to a situation in which the value of the bubble exceeds the wealth of the young who must buy it. Tirole (1985) added population growth, finding that coexis-

tence is possible when the rate of return on investment does not exceed the rate of population growth, so that the per-capita value of the bubble remains bounded. We note that we do not know whether coexistence is possible in a stochastic economy without population growth, either with or without ergodicity. For a model of stationary monetary equilibrium, see Gottardi (1991).

3. MARKET EQUILIBRIUM WITH HETEROGENEOUS LONG-LIVED AGENTS

3.1. *Introduction*

The model presented by R. Lucas in "Asset Prices in an Exchange Economy" (1978) has had an important influence on financial economics, macroeconomics, monetary theory, and econometrics. As in the precursor by LeRoy (1973), the setting of the model is made up of a time-homogeneous Markov shock process, a single infinitely-lived agent with additive discounted utility for consumption of a single commodity, and markets for securities paying dividends in each period specified by a fixed function of the shock in that period. Prescott and Mehra (1980) extended the model to the case of a class of agents with identical preferences and identical endowments. For purposes of equilibrium, one still has the effect of a single representative agent. With a sufficient number and type of securities to create the effect of complete markets, heterogeneity of agents can be accommodated because the existence of an equilibrium follows as a straightforward application of the results of Bewley (1972) (see Levine and Zame (1992) and Magill and Quinzii (1992)). By virtue of the first welfare theorem, any such equilibrium is the equilibrium of an appropriately chosen representative agent, so it will display the properties derived in Lucas' analysis. With incomplete markets and heterogeneous agents, however, it is well known that equilibrium allocations are not typically efficient, ruling out the possibility of a representative agent. This section will be devoted to establishing conditions under which one can still recover the existence of a stationary Markov equilibrium with an ergodic measure in a Lucas model with heterogeneous agents *and* incomplete markets.

At the equilibrium of Lucas' original results, the single agent must hold the entire market portfolio and consume all dividends in each period. With additively separable and stationary preferences, this implies that prices are given by time-invariant functions of just the current shock. With heterogeneous agents, however, we cannot hope for so much: equilibrium prices must clearly depend at least on the distribution of assets across the various agents. For example, with a fixed current shock, reallocation of the securities from one agent to another with greater risk aversion would typically require a new set of equilibrium prices. (In fact, even if the agents have identical preferences but begin with different endowments, equilibrium prices depend on the particular security allocation, which changes through time.) Thus, in order to describe equilibrium in terms of a single time-homogeneous Markov process, one must include both endogenous and exogenous variables in the state description.

3.2. *The Primitives*

There is a finite set Y of *exogenous states* or *shocks*. The random process $\{y_t\}$ of exogenous states is a time-homogeneous Markov process with transition $P: Y \rightarrow \mathcal{P}(Y)$. The set of dates for the economy is either $\{1, \dots, T\}$ for some finite T , or \mathbb{N} , in which case we will write “ $T = +\infty$.”

There are n long-lived securities and m long-lived agents. (One should actually imagine that there are m agent types, where for each type there is a continuum of identical agents; as usual, speaking of m individuals eases the discussion.) For notational ease, consumption dividends of the n available securities are defined by $d: Y \rightarrow \mathbb{R}_+^n$, where $d_j(y)$ is the dividend paid by security j when the current shock is y . In addition, $e: Y \rightarrow \mathbb{R}^m$ defines the private consumption endowments of the m agents; that is, agent i receives $e_i(y)$ when the current shock is y .

The von Neumann-Morgenstern preferences of agent i are defined by a strictly monotone C^1 concave function $u_i: \mathbb{R}_{++} \rightarrow \mathbb{R}$ that is bounded above and unbounded below, and a discount factor $\beta_i \in (0, 1)$. For any sequence $C = \{C_t\}_{t=1}^T$ of \mathbb{R}_{++} -valued random variables on some probability space, the associated utility for agent i is

$$U_i^T(C) = E \left[\sum_{t=1}^T \beta_i^t u_i(C_t) \right].$$

Note that, since u_i is unbounded below, there is some constant $\underline{c}_i > 0$ such that

$$u_i(\underline{c}_i) + \frac{\beta_i}{1 - \beta_i} \bar{u}_i \leq \frac{1}{1 - \beta_i} u_i \left(\min_{y \in Y} e_i(y) \right),$$

where \bar{u}_i is an upper bound on u_i . It is “irrational” for agent i to have consumption below \underline{c}_i in any period, since such behavior is dominated by consuming endowments then and in all subsequent periods.

With $u = (u_1, \dots, u_m)$ and $\beta = (\beta_1, \dots, \beta_m)$, the primitives of the economy are $\mathcal{E} = (P, e, u, \beta, d)$. We mention at this point that it would be possible, at the expense of greater notational complexity, to generalize this structure in several ways without changing the essential features of the analysis. In particular, we could allow the number of consumption goods in each period to be any finite number, and we could allow the felicity functions u_i to depend on the shock; for details see an earlier draft of this paper.

Let $\Lambda = \{\theta \in (\mathbb{R}_+^n)^m: \sum_{i=1}^m \theta_i = (1, \dots, 1)\}$ be the set of possible vectors of portfolios at a point in time. Note that we do not allow short sales. Let Λ^- be a copy of Λ . The endogenous state space for the economy is $Z = \Lambda^- \times \Lambda \times \mathbb{R}_{++}^n \times \mathbb{R}_+^n$, with typical element (θ^-, θ, c, p) representing the vector p of security prices and, for each agent i , the portfolio $\theta_i^- \in \mathbb{R}_+^n$ of securities purchased in the previous period, the portfolio θ_i purchased in the current period, and the consumption level c_i .

The state space for the economy (with feasibility “built in”) is

$$S = \left\{ [y, (\theta^-, \theta, c, p)] \in Y \times Z \mid \sum_i [c_i - e_i(y)] = \sum_j d_j(y) \right\}.$$

For an S -valued stochastic process $\{s_t\}_{t=1}^T$ on some probability space, we typically write $s_t = [y_t, (\theta_t^-, \theta_t, c_t, p_t)]$. We call such a process $\{s_t\}$ a *consistent state process* for \mathcal{E} provided that, for all $t < T$, the conditional distribution of y_{t+1} given $\{s_1, \dots, s_t\}$ is almost surely $P(y_t)$. A sequence $\{a_t\}$ of random variables (on the same probability space on which $\{s_t\}$ is defined) is *adapted* to $\{s_t\}$ if a_t is (s_1, \dots, s_t) -measurable for all t .

Given a state process $\{s_t\}$, a *feasible policy* for agent i is an $\{s_t\}$ -adapted process $\{(\theta_i, C_i)\}$ valued in $\mathbb{R}_+^n \times \mathbb{R}_{++}$ and satisfying, for all t ,

$$p_t \cdot \theta_i + C_i = \theta_{i,t-1} \cdot [p_t + d(y_t)] + e_i(y_t),$$

with $\theta_{i,0}$ defined to be $\theta_{i,1}^-$. Likewise, an *optimal policy* for i is a feasible policy $\{(\theta_i, C_i)\}$ such that, for any other feasible policy $\{(\theta'_i, C'_i)\}$, we have $U_i^T((C_i)) \geq U_i^T((C'_i))$.

DEFINITION 3.1: An *equilibrium* for \mathcal{E} is a consistent state process $\{s_t\}$ with the property that, for all i , the policy $\{(\theta_i, c_{it})\}$ is optimal given $\{s_t\}$.

This is the standard Walrasian equilibrium concept. It is encompassed in our framework as follows.

DEFINITION 3.2: An *equilibrium transition* for \mathcal{E} is a pair (J, Π) , where J is a measurable subset of S and $\Pi: J \rightarrow \mathcal{P}(J)$ is a transition with the properties:

- (a) for any $(\bar{\theta}, \bar{y}) \in \Lambda^- \times Y$, there is a point $[\bar{y}, (\bar{\theta}, \theta, c, p)] \in J$ (that is, any starting parameters are admissible);
- (b) each time-homogeneous Markov J -valued process $\{s_t\}$ with transition Π is an equilibrium for \mathcal{E} with $T = +\infty$.

The main result of this section is:

THEOREM 3.1: *There exists an equilibrium transition (J, Π) for \mathcal{E} having an ergodic measure.*

As we remark at the end of this section, our method of proof actually implies the existence of conditionally spotless equilibrium, in the sense described in Section 1.6. The proof of the theorem proceeds in two basic steps.

3.3. First Step: The Expectations Correspondence

First, we define a set-valued function $g: S \rightarrow \mathcal{P}(S)$ by letting $\mu \in g(s_0)$, where $s_0 = [y_0, (\theta_0^-, \theta_0, c_0, p_0)]$, if:

- (i) the support of μ is the graph of some function $h: Y \rightarrow Z$;
- (ii) the marginal μ_Y of μ on Y is $P(y_0)$ and the marginal of μ on Λ^- is δ_{θ_0} almost surely;
- (iii) for all i , $p_0 \cdot \theta_{i0} + c_{i0} = \theta_{i0}^- \cdot [p_0 + d(y_0)] + e_i(y_0)$;
- (iv) for all i , $\lambda_i \equiv \beta_i E(u'(c_{i1})[p_1 + d(y_1)]) - u'(c_{i0})p \geq 0$ and $\lambda_i \cdot \theta_{i0} = 0$, for any random variable $s_1 = [y_1, (\theta_1^-, \theta_1, c_1, p_1)]$ with distribution μ ;
- (v) $c_1 \geq \underline{c}$ almost surely, where $\underline{c} = (\underline{c}_1, \dots, \underline{c}_m)$ is as defined in Section 3.2.

PROPOSITION 3.1: *There exists a compact self-justified set J for g with the property: For any $(\bar{\theta}, \bar{y}) \in \Lambda^- \times Y$, there exists $[\bar{y}, (\bar{\theta}, \theta, c, p)] \in J$.*

The proposition is proved by establishing the following five lemmas. For any time horizon T , let $\eta(T)$ denote the corresponding set of equilibria for \mathcal{E} .

The first lemma follows from the property of \underline{c}_i described in Section 3.2.

LEMMA 3.1: *For all T , if $\{s_t\} \in \eta(T)$, then for all i and t , $c_{it} \geq \underline{c}_i$ almost surely.*

We next guarantee a uniform upper bound on finite horizon equilibrium security prices.

LEMMA 3.2: *There exists $\bar{p} \in \mathbb{R}_+^n$ such that, for any finite T , if $\{s_t\} \in \eta(T)$, then $p_t \leq \bar{p}$ for all t almost surely.*

PROOF: Let $\{(\theta_{it}, c_{it})\}$ denote the policy of agent i in equilibrium. Consider the following modification of a policy: In the current period, replace the final portfolio θ_{it} with $(1 - \varepsilon)\theta_{it}$, spending $\varepsilon p_t \cdot \theta_{it}$ on additional current consumption, and replace the pairs $(\theta_{i\tau}, c_{i\tau})$, $\tau > t$, with $(1 - \varepsilon)(\theta_{i\tau}, c_{i\tau})$. Clearly the modified policy satisfies the budget constraints of all periods.

This modification leads to a loss of expected utility in future periods that is bounded, for small ε by $K_i \varepsilon$, for some K_i that is independent of T and the particular equilibrium under consideration. [One can take $K_i = 1 + \sup\{cu'_i(c) : \underline{c}_i \leq c \leq \sum_k e_k(y) + \sum_j d_j(y)\}$.]

If asset prices can be arbitrarily high, then the value of the economy's wealth, and thus the value of the richest agent's portfolio, can also be arbitrarily large. Consequently, the fraction of this agent's portfolio that must be sold in order to finance one unit of consumption, say, can be arbitrarily small. In equilibrium, the incremental utility generated by adding one unit to this agent's consumption is bounded below (this follows easily from Lemma 3.1). Thus, if asset prices could be arbitrarily large, there would be the possibility, for some agent, of obtaining this increment of utility while giving up an arbitrarily small amount of future utility, a contradiction of optimality. Q.E.D.

LEMMA 3.3: For any $T \geq 2$, if $\{s_t\}_{t=1}^T \in \eta(T)$, then $\{s_t\}_{t=2}^T \in \eta(T-1)$.

PROOF: This follows from the Bellman's principle of optimality. *Q.E.D.*

For any T , let $S_T = \{s \in S \mid \text{there exists } \{s_t\} \in \eta(T+1) \text{ with } s_1 = s\}$, the set of initial equilibrium states in some T -horizon equilibrium.

LEMMA 3.4: For all finite T and all $(\bar{\theta}, \bar{y}) \in \Lambda^- \times Y$, there exists $[\bar{y}, (\bar{\theta}, \theta, c, p)] \in S_T$.

PROOF: Radner's (1972) existence result applies when security prices can be shown to be strictly positive in any quasi-equilibrium. Since we don't require strictly positive dividends, one must extend Radner's proof. A suitable proof can be made available on request. See Grodal and Vind (1988) for an alternative proof. *Q.E.D.*

Let $K = \{[y, (\theta^-, \theta, c, p)] \in S \mid c \geq \underline{c}, p \leq \bar{p}\}$, where \bar{p} is given by Lemma 3.2, and let C_{0j} be defined as in Section 1.4.

LEMMA 3.5: For all T , $S_T \subset C_{0T}$.

PROOF: The proof is by an easy induction argument, using Lemma 3.3. *Q.E.D.*

The compactness of K and Lemmas 3.4 and 3.5, combined with Theorem 1.2 of the central results, together prove Proposition 3.1.

3.4. Second Step: Optimality of Candidate Policies

Let $G: S \rightarrow \mathcal{P}(S)$ be defined by letting $G(s)$ be the closure of the convex hull of $g(s)$. We do not assert that G is an expectations correspondence, in the sense of having a closed graph. Nonetheless, by the previous proposition and the central results (Proposition 1.3), there is an ergodic Markov equilibrium (J, Π, μ) for G . In particular, for any $(\bar{\theta}, \bar{y}) \in \Lambda^- \times Y$, there is a time-homogeneous J -valued Markov process $\{s_t\}$ transition function Π with $(\theta_1^-, y_1) = (\bar{\theta}, \bar{y})$. Theorem 3.1 is thus implied by the following proposition.

PROPOSITION 3.2: Given a state process $\{s_t\} = \{[y_t, (\theta_t^-, \theta_t, c_t, p_t)]\}$ with transition function Π , for any agent $i = 1, \dots, m$, the corresponding policy $\{(\theta_{it}, c_{it})\}$ is optimal.

PROOF: Given $\{s_t\}$ with transition Π , let $\{(\theta_t, C_t)\}$ be an arbitrary budget-feasible policy for i . We will show that $U_i^\infty(\{c_{it}\}) \geq U_i^\infty(\{C_t\})$, which proves the optimality of $\{(\theta_{it}, c_{it})\}$.

Let T be arbitrary. We will first show that

$$(3.1) \quad U_i^\infty(\{c_{it}\}) \geq \dot{U}_i^T(\{C_i\}) + E \left[\sum_{t=T+1}^{\infty} \beta_i^t u_i(c_{it}) \right] \\ + \beta_i^T E[u_i'(c_{iT}) p \cdot (\Theta_T - \theta_{iT})].$$

We can prove (3.1) by induction. It is certainly true for $T = 1$ since the concavity of u_i implies that $u_i(c_{i1}) \geq u_i(C_1) + u_i'(c_{i1})(c_{i1} - C_1)$. Assuming that (3.1) holds for $T = \tau$, we then show that it holds for $T = \tau + 1$. The definition of G implies that

$$u_i'(c_{i\tau}) p_\tau \cdot (\Theta_\tau - \theta_{i\tau}) \geq \beta_i E(u_i'(c_{i,\tau+1}) [p_{\tau+1} + d(y_{\tau+1})] | s_1, \dots, s_\tau) \\ \cdot (\Theta_\tau - \theta_{i\tau}) \quad \text{a.s.}$$

We also have, again by concavity,

$$u_i(c_{i,\tau+1}) \geq u_i(C_{\tau+1}) + u_i'(c_{i,\tau+1})(c_{i,\tau+1} - C_{\tau+1}).$$

Then (3.1) follows for $T = \tau + 1$ by combining the last two expressions with the law of iterated expectations and the fact that

$$[p_{\tau+1} + d(y_{\tau+1})] \cdot (\Theta_\tau - \theta_{i\tau}) + c_{i,\tau+1} - C_{\tau+1} = p_{\tau+1} \cdot (\Theta_{\tau+1} - \theta_{i,\tau+1}).$$

Thus (3.1) follows for all T by induction. Since u_i is bounded above and $c_{it} \geq c_i$ a.s., the second term on the right-hand side of (3.1) converges to zero as $T \rightarrow +\infty$. Since $c_{it} \geq \underline{c}_i$ a.s. and u_i is concave, $0 \leq u_i'(c_{it}) \leq u_i'(\underline{c}_i)$ a.s., so the third term on the right-hand side of (3.1) also converges to zero with T , implying the optimality of $\{c_{it}\}$. Q.E.D.

This completes the proof of Theorem 3.1.

3.5. Equilibrium Properties and Concluding Remarks

As in the previous section, our analysis allows the distribution of states in a given period to depend on more than a minimal number of prior variables. For example, the definition of the expectations correspondence used to prove Theorem 3.1 has no conditions in which the specific portfolio an agent holds at the beginning of a period has any role that is not summarized by the agent's initial wealth. As shown in Section 1.5, there exist equilibrium transitions in which the distribution of states in the next period depends only on a suitably defined "minimal summary" of the current state.

Since the space Y of shocks is finite, the equilibria demonstrated by our method of proof are conditionally spotless, in the sense of Section 1.6.

Since these results were first presented, they have been used as the basis for applied work by Marcet and Singleton (1991) and have been extended by Ma (1993) to the case of recursive utility, in the sense of Epstein and Zin (1989). Under a concavity assumption on the recursive utility aggregator, an expectations correspondence can once again be used to characterize equilibrium in terms of joint conditions on current state variables and the distribution of state

variables in the following period. By the nature of nonadditive recursive utility, however, this entails an enlargement of the state space to include continuation utilities, as in Section 4.

Without assuming the short-sales constraint that we exploit here, Magill and Quinzii (1992), Levine and Zame (1992), and Hernandez and Santos (1991, 1992) have shown the existence of (not necessarily Markov or ergodic) equilibrium in incomplete markets, under various sets of assumptions (otherwise, see Hart's (1975) counterexample) about the structure of security markets and budget constraints.

4. STOCHASTIC GAMES

4.1. *Introduction*

In this section we study what our central results yield for nonzero sum stochastic games, as defined by Shapley (1953). This is a quite general model of strategic interaction with symmetric information; repeated games, for example, are the special case given by a singleton state space. For finite, or countable, state spaces there is a variety of existence theorems for time homogeneous Markovian Nash equilibrium. See Shapley (1953), Parthasarathy (1982), Nowak (1985), or Rieder (1979). Although existence of stationary Markovian Nash equilibrium has been established in several special cases with uncountable state space (Himmelberg et al. (1976) and Parthasarathy (1982) for transitions satisfying certain separability conditions, Parthasarathy and Sinha (1989) for state-independent transitions, Majumdar and Sundaram (1991) for state and action spaces that are compact intervals, with symmetric payoffs), and there are ε -equilibrium results (Whitt (1980), Nowak (1985)), a general treatment has been lacking. For recent related work, see Bergin and Bernhardt (1991), Nowak and Raghavan (1988), Nowak (1992), and Börgers (1991).

Here we apply our central results to the finite-horizon equilibrium existence theorem of Rieder (1979). We do not obtain the existence of a stationary Markovian equilibrium. But we do get the existence of a special kind of correlated equilibrium in which (within the conventional description of correlated equilibrium) the messages transmitted by the mediator are public. In a one-stage game, correlated equilibrium with publicly observed messages is equivalent to randomization over the set of Nash equilibria, and a similar interpretation is appropriate here. In formulating myopic equilibrium conditions we expand the state space, including a proxy for the expected utilities of the continuation of the game. (The ideas used in Sections 2 and 3 to derive true optimization from myopic optimization are not applicable here.) These components of the expanded states serve as signalling devices for selecting continuation equilibria. If we were to allow dependence on complete past histories, we still could not avoid a minimal use of "sunspot" variables if our goal was to have an ergodic measure.

Mertens and Parthasarathy (1987) have established the existence of a (subgame perfect) Nash equilibrium for stochastic games with general state and

action sets under hypotheses of the same nature, but somewhat weaker than ours. (See also Amir (1987).) Our researches have been independent. Theirs is a remarkable result, with which ours is, in comparison, elementary. There is, however, an important reason why our result is not covered by theirs, namely, that we can guarantee that the equilibrium process admits an ergodic measure.

Nowak and Raghavan (1988), also working independently, establish existence of a stationary correlated equilibrium under more general assumptions than those considered here. Due to the special nature of the correlation employed here (publicly observed signals) and the fact that we obtain an ergodic measure, our result is not a special case of theirs. Nowak and Raghavan (1992) obtain a result quite similar to ours, but again without the existence of an ergodic measure, using somewhat weaker assumptions. Nowak (1992) extends the analysis to equilibria with limit-of-means payoffs.

In a setting of continuous games, Harris, Reny, and Robson (1993) make a case for correlation in extensive form games based on public signals, using some techniques related to those in this paper. In their case, however, the issue is existence of any subgame perfect equilibria, which need not be in any sense stationary, and their use of correlation is to allow the informational significance of distinctions between actions to be preserved in the limit as the distinction itself disappears.

4.2. Stochastic Games and Extensions

An n -player stochastic game is a tuple $\Gamma = (Y, A, P, u, \beta)$, where:

- (i) Y is a nonempty compact metric space of *shocks*;
- (ii) $A = A_1 \times \cdots \times A_n$, where each A_i is a compact metric space of *actions*;
- (iii) $P: Y \times A \rightarrow \mathcal{P}(Y)$ is a transition probability, called the *law of motion*;
- (iv) $u = (u_1, \dots, u_n)$ is an n -tuple of continuous short-run *utility functions*,

where $u_i: Y \times A \rightarrow \mathbb{R}$;

- (v) $\beta = (\beta_1, \dots, \beta_n) \in (0, 1)^n$ is an n -tuple of *discount factors*.

We fix a game $\Gamma = (Y, A, P, u, \beta)$ for the remainder.

Intuitively, an “extension” of Γ involves an expanded state space containing public information that does not affect the utilities or transitions on the original state space, but which can serve as a source of signals for coordinating behavior. Formally, $\Gamma^* = (Y^*, A, P^*, u^*, \beta)$ is an extension of Γ if there is a continuous surjection $\pi: Y^* \rightarrow Y$ (the “projection”) such that, for all $(y^*, a) \in Y^* \times A$:

- (i) $P^*(y^*, a) \circ \pi^{-1} = P(\pi(y^*), a)$;
- (ii) for all i , $u_i^*(y^*, a) = u_i(\pi(y^*), a)$.

4.3. Strategies and Equilibrium

Typical elements of $\mathcal{P}(A_i)$ are denoted by α_i . Whenever convenient we identify $\prod_i \mathcal{P}(A_i)$ with the obvious subset of $\mathcal{P}(A)$; that is, we identify each $\alpha = (\alpha_1, \dots, \alpha_n)$ with $\alpha_1 \times \cdots \times \alpha_n \in \mathcal{P}(A)$.

A (stationary, 0-th order, mixed) *strategy* for player i in Γ is a transition probability $\xi_i: Y \rightarrow \mathcal{P}(A_i)$. A strategy vector is then some $\xi = (\xi_1, \dots, \xi_n)$, a function from Y to $\prod_i \mathcal{P}(A_i)$.

The interpretation of Γ as a game results from the specification of objective functions for the agents. Given a strategy vector ξ and an *initial distribution* $\sigma \in \mathcal{P}(Y)$, the i th player's expected utility is

$$U_i(\xi; \sigma) = E \left[\sum_{t=0}^{\infty} \beta_t^i u_i(y_t, a_t) \right],$$

where the expectation is taken with respect to the distribution of the time-homogeneous Markov process $\{(y_t, a_t)\}$ that is generated by σ and the transition probabilities P and ξ . Specifically, the joint distribution of $(y_0, a_0, y_1, a_1, \dots)$ is $\sigma \otimes \xi \otimes P \otimes \xi \otimes P \otimes \dots$. We will also have occasion to refer to the transition function $Q_\xi: Y \rightarrow \mathcal{P}(Y)$ for the shock process given by $Q_\xi(y) = [(\delta_y \times \xi(y)) \otimes P]_Y = \int_A P(y, a) \xi(y; da)$.

A strategy vector $\bar{\xi}$ is a σ -*equilibrium* if $U(\bar{\xi}; \sigma) \geq U_i(\bar{\xi}_{-i}, \xi_i; \sigma)$ for all ξ_i and i . (Here, $(\bar{\xi}_{-i}, \xi_i)$ stands for the strategy $\bar{\xi}$, except that the i th coordinate strategy has been replaced by ξ_i .) A strategy vector $\bar{\xi}$ is an *equilibrium* if it is a σ -equilibrium for any initial distribution σ . Clearly, $\bar{\xi}$ is an equilibrium if and only if it is a δ_y -equilibrium for all y in Y . If none of the agents is conditioning behavior on the history of play, or other publicly available information not encoded in the state, then no single agent can derive any advantage from doing so, but it can also happen that such information is used to coordinate behavior.

DEFINITION 4.1: A *coordinated equilibrium* for a game Γ is an equilibrium ξ^* for some extension Γ^* of Γ . A *coordinated equilibrium* ξ^* admits an *ergodic measure* if there is an ergodic measure for the associated transition $Q_{\xi^*}: Y^* \rightarrow \mathcal{P}(Y^*)$.

Our existence theorem employs the following continuity hypotheses on the transition probability $P: Y \times A \rightarrow \mathcal{P}(Y)$:

- (i) (bounded variation) norm continuity, that is, if $(y_n, a_n) \rightarrow (y, a)$, then the supremum of $|P(y_n, a_n)(B) - P(y, a)(B)|$ over all Borel sets B goes to zero, and
- (ii) any measure $P(y, a)$ is absolutely continuous with respect to any other $P(y', a')$; that is, $P(y, a) \ll P(y', a')$.

THEOREM 4.1: Any game Γ with norm-continuous and mutually absolutely continuous transition probabilities has a coordinated equilibrium admitting an ergodic measure.

4.4. The Expectations Correspondence

We apply the central results to the following finite-stage equilibrium existence theorem.

PROPOSITION 4.1 (Rieder (1979)): *Suppose that for each $y \in Y$, $P(y, \cdot): A \rightarrow \mathcal{P}(Y)$ is weakly continuous. Then for any integer $T \geq 0$ and any universally measurable functions $v_{1,T+1}, \dots, v_{n,T+1}: Y \rightarrow \mathbb{R}$ there exist transition functions $\bar{\xi}_0 = (\bar{\xi}_{1,0}, \dots, \bar{\xi}_{n,0}), \dots, \bar{\xi}_T = (\bar{\xi}_{1,T}, \dots, \bar{\xi}_{n,T})$ such that for any initial distribution $\sigma \in \mathcal{P}(Y)$, any i , and any transition functions $\xi_{i,0}, \dots, \xi_{i,T}: Y \rightarrow \mathcal{P}(A_i)$,*

$$\begin{aligned} & \int \left(\sum_{t=0}^T \beta_t^i u_i(y_t, a_t) + v_{i,T+1}(y_{T+1}) \right) d[\sigma \otimes \bar{\xi}_0 \otimes P \otimes \dots \otimes \bar{\xi}_T \otimes P] \\ & \geq \int \left(\sum_{t=0}^T \beta_t^i u_i(y_t, a_t) + v_{i,T+1}(y_{T+1}) \right) \\ & \quad \times d[\sigma \otimes (\bar{\xi}_{-i,0}, \xi_{i,0}) \otimes P \otimes \dots \otimes (\bar{\xi}_{-i,T}, \xi_{i,T}) \otimes P]. \end{aligned}$$

For the purposes of applying the central results, our state space shall be $S = Y \times Z$, where

$$Z = \prod_i \mathcal{P}(A_i) \times [-r, r]^n,$$

with r chosen so that, for all i ,

$$u_i(Y \times A) \subset (1 - \beta_i)[-r, r].$$

We let $y(s)$, $\alpha(s)$, and $v(s)$ denote the projections of a generic s onto Y , $\prod_i \mathcal{P}(A_i)$, and $[-r, r]^n$ respectively, and we sometimes write " $s = (y, \alpha, v)$."

For a given $\bar{s} = (\bar{y}, \bar{\alpha}, \bar{v})$, $G(\bar{s}) \subset \mathcal{P}(S)$ is the set formed by those $\nu \in \mathcal{P}(S)$ having the three properties: (i) ν is compatible with the transition probabilities implied by \bar{y} and $\bar{\alpha}$, (ii) \bar{v} equals the vector of discounted sums of expected utilities implied by $\bar{\alpha}$ and the transition probabilities, and (iii) no player i has an incentive to replace $\bar{\alpha}_i$ by another α_i if the change leaves unaltered the family of probabilities on Z conditional on Y . The last condition requires some explanation. Let $K = P(Y \times A)$. Since P is norm-continuous, K is norm-compact, and the assumption of mutual absolute continuity assures that $\rho \ll \rho'$ for any ρ and ρ' in K . Denote $\mathcal{P}_K(Y \times Z) = \{\nu \in \mathcal{P}(Y \times Z) | \nu_Y \in K\}$. For every $\nu \in \mathcal{P}(Y \times Z)$, let $\nu^Y: Y \rightarrow \mathcal{P}(Z)$ be a version of regular conditional probability for ν , that is, $\nu_Y \otimes \nu^Y = \nu$. Define $\Psi: K \times \mathcal{P}_K(Y \times Z) \rightarrow \mathcal{P}_K(Y \times Z)$ by $\Psi(\rho, \nu) = \rho \otimes \nu^Y$. (Since $\rho \ll \nu_Y$, this definition is independent of the version of conditional probability chosen.) In words, the operator Ψ combines ρ and ν into a new probability measure on $Y \times Z$ having ρ as marginal on Y , but with the probabilities on Z conditional on Y remaining the same as those in ν .

We may now state, precisely, that $\nu \in G(\bar{s})$ if and only if, for all i ,

- (i) $\nu_Y = \int_A P(\bar{y}, \cdot) d\bar{\alpha}$,
- (ii) $\bar{v}_i = \int_A u_i(\bar{y}, \cdot) d\bar{\alpha} + \beta_i \int_S v_i d\nu$, and
- (iii) $\bar{v}_i \geq \int_A u_i(\bar{y}, \cdot) d(\bar{\alpha}_{-i}, \alpha_i) + \beta_i \int_S v_i d\Psi[\int_A P(\bar{y}, \cdot) d(\bar{\alpha}_{-i}, \alpha_i), \nu]$ for all $\alpha_i \in \mathcal{P}(A_i)$.

We prove the following in Section 4.7.

PROPOSITION 4.2: $G: S \rightarrow \mathcal{P}(S)$ is convex-valued with a closed graph.

4.5. Ergodic Markov Equilibrium for G

PROPOSITION 4.3: The expectations correspondence G has an ergodic Markov equilibrium.

PROOF: We verify that, in the notation of Section 1, subsection 1.4, $C_{0,T+1} \neq \emptyset$ for all $T = 1, 2, \dots$. Since S is compact, Theorem 1.2 then yields a compact self-justified set, and the existence of an ergodic Markov equilibrium follows from Corollary 1.1.

Fixing a particular T , let $\bar{\xi}_0 = (\bar{\xi}_{1,0}, \dots, \bar{\xi}_{n,0}), \dots, \bar{\xi}_T = (\bar{\xi}_{1,T}, \dots, \bar{\xi}_{n,T})$ satisfy the conclusions of Proposition 4.1 when the terminal reward functions $v_{1,T+1}, \dots, v_{n,T+1}$ are identically zero. For $t = 0, \dots, T$ and $i = 1, \dots, n$ define $w_{it}: Y \rightarrow \mathbb{R}$ by

$$w_{it}(y) = \int \left(\sum_{\tau=t}^T \beta_i^{\tau-t} u_i(y_\tau, a_\tau) \right) d[\delta_y \otimes \bar{\xi}_t \otimes P \otimes \dots \otimes P \otimes \bar{\xi}_T].$$

For each $y \in Y$ and $t = 0, \dots, T$ set

$$\bar{s}_t(y) = (\bar{y}_t(y), \bar{\alpha}_t(y), \bar{v}_t(y)) = (y, \bar{\xi}_t(y), w_{it}(y)).$$

Choosing $\alpha_{T+1} = (\alpha_{1,T+1}, \dots, \alpha_{n,T+1}) \in \prod_i \mathcal{P}(A_i)$ arbitrarily, for $y \in Y$ set

$$\bar{s}_{T+1}(y) = (\bar{y}_{T+1}(y), \bar{\alpha}_{T+1}(y), \bar{v}_{T+1}(y)) = (y, \alpha_{T+1}, 0).$$

For each $y \in Y$ and $t = 0, \dots, T$, define $\nu_t(y) \in \mathcal{P}(S)$ by

$$\nu_t(y)(E) = Q_{\bar{\xi}_t}(y)(\{y' \in Y \mid \bar{s}_{t+1}(y') \in E\}).$$

We claim that $\nu_t(y) \in G(\bar{s}_t(y))$ for $[\sigma \otimes \bar{\xi}_0 \otimes P \otimes \dots \otimes P \otimes \bar{\xi}_{t-1} \otimes P]$ -a.e. $y \in Y$. Indeed, conditions (i) and (ii) above hold automatically, and (iii) follows from the maximization assertion of Proposition 4.1 once one recognizes that the integrand of the second integral on the right-hand side of (iii) depends only on the shock. Let $C_{00} = \bigcup_{t=0}^{T+1} \bar{s}_t(Y)$. Then for every $y \in Y$ we have $\nu_T(y)(C_{00}) = 1$, so that $\nu_T(y)(C_{00}) = 1$ and $\bar{s}_T(y) \in C_{00}$. In general, if $\bar{s}_t(y) \in C_{0,T+1-t}$ for $[(\sigma \otimes \bar{\xi}_0 \otimes P \otimes \dots \otimes P \otimes \bar{\xi}_{t-1} \otimes P) \circ \bar{s}_t^{-1}]$ -a.e. $s \in S$, then $\nu_{t-1}(y)(C_{0,T+1-t}) = 1$ for $[(\sigma \otimes \bar{\xi}_0 \otimes P \otimes \dots \otimes P \otimes \bar{\xi}_{t-2} \otimes P)]$ -a.e. $y \in Y$, whence $\bar{s}_{t-1}(y) \in C_{0,T+1-(t-1)}$ for $(\sigma \otimes \bar{\xi}_0 \otimes P \otimes \dots \otimes P \otimes \bar{\xi}_{t-2} \otimes P) \circ \bar{s}_{t-1}^{-1}$ -a.e. $s \in S$. By induction it follows that $\bar{s}_0(y) \in C_{0,T+1}$ for $(\sigma \circ \bar{s}_0^{-1})$ -a.e. $s \in S$, and in particular $C_{0,T+1} \neq \emptyset$, as desired. Q.E.D.

4.6. Defining an Equilibrium of an Extension Γ^* of the Game Γ

Theorem 4.1 is now proved by transforming (J, Π, μ) into a coordinated equilibrium for the game Γ .

Let $Y^* = J$, and define $\pi: Y^* \rightarrow Y$, $u_i^*: Y^* \times A \rightarrow \mathbb{R}$, and $P^*: Y^* \times A \rightarrow \mathcal{P}(Y^*)$ by:

$$\begin{aligned} \pi(y, \alpha, \nu) &= y; & u_i^*(y^*, a) &= u_i(\pi(y^*), a); \\ P^*(y^*, a) &= \Psi(P(\pi(y^*), a), \Pi(y^*)). \end{aligned}$$

Consequently $P^*(y^*, a) \circ \pi^{-1} = P(\pi(y^*), a)$, and the following result, proved in Section 4.7, completes the verification that $\Gamma^* = (Y^*, A, P^*, u^*, \beta)$ is an extension of Γ .

PROPOSITION 4.4: $P^*: Y^* \times A \rightarrow \mathcal{P}(Y^*)$ is measurable.

Using the notation $y^* = (y(y^*), \alpha(y^*), \nu(y^*))$, we define the strategy $\xi^*: Y^* \rightarrow \prod_i \mathcal{P}(A_i)$ simply by letting $\xi^*(y^*)$ be the projection of y^* on the α coordinate; that is, $\xi^*(y^*) = \alpha(y^*)$. Since the transition $\Pi = Q_{\xi^*}$ has an ergodic measure μ by Proposition 4.3, the next proposition concludes the proof of Theorem 4.1.

PROPOSITION 4.5: The strategy vector ξ^* constitutes an equilibrium for Γ^* .

PROOF: This is a straightforward dynamic programming argument.

Consider player i . Because utility functions are bounded and $\beta_i < 1$, conditions (ii) and (iii) of the definition of the expectations correspondence G yield, first, $U_i(\xi^*; \delta_{y^*}) = v_i(y^*)$ for all y^* and, second,

$$v_i(y^*) \geq \int_A \left[u_i^*(y^*, a) + \beta_i \int_{Y^*} v_i(\cdot) dP^*(y^*, a) \right] d[\xi_{-i}^*(y^*), \delta_{a_i}](a)$$

for all y^* and a'_i . Of course, this says that v_i satisfies the Bellman equation for the dynamic programming problem facing player i . Since the policy ξ_i^* realizes v_i , it follows that ξ_i^* is optimal against ξ_{-i}^* for any initial distribution σ . This completes the proof of Theorem 4.1 with the exception of the technical arguments to follow. Q.E.D.

4.7. Technical Arguments

For clarity we distinguish between $\mathcal{P}(Y)$, the set of probability measures on Y with the weak convergence topology, and $\hat{\mathcal{P}}(Y)$, the same set but with the norm topology.

LEMMA 4.1: The function $\Psi: K \times \mathcal{P}_K(Y \times Z) \rightarrow \mathcal{P}_K(Y \times Z)$ is continuous (with the norm topology on K).

PROOF: We first show that if $\nu_n \rightarrow \nu$ weakly, with $\nu_n \in \mathcal{P}_K(Y \times Z)$, then $\nu_{nY} \rightarrow \nu_Y$ in norm. Indeed, the collection $\{\nu_{nY}\}$ is a subset of K and so has at least one strong, hence weak, accumulation point ρ . However, $\nu_{nY} \rightarrow \nu_Y$ weakly,

so ν_Y is the only weak accumulation point. Therefore $\rho = \nu_Y$ and we conclude that ν_Y is also the only strong accumulation point, that is, $\nu_{nY} \rightarrow \nu_Y$.

Let now $\nu_n \rightarrow \nu$ weakly and $\rho_n \rightarrow \rho$ in norm. Since

$$\Psi(\rho_n, \nu_n) = \rho_n \otimes \nu_n^Y = (\rho_n - \rho) \otimes \nu_n^Y + \rho \otimes \nu_n^Y$$

and $(\rho_n - \rho) \otimes \nu_n^Y \rightarrow 0$, it suffices to show that $\rho \otimes (\nu_n^Y - \nu^Y) \rightarrow 0$.

Observe that $\nu_Y \otimes (\nu_n^Y - \nu^Y) = \nu_n - \nu - (\nu_{nY} - \nu_Y) \otimes \nu_n^Y$. Because $\nu_n \rightarrow \nu$ weakly and $\nu_{nY} \rightarrow \nu_Y$ in norm, we have $\nu_Y \otimes (\nu_n^Y - \nu^Y) \rightarrow 0$ weakly.

Since $\rho \ll \nu_Y$, there is a (Radon-Nikodym) derivative $g \in L^1(Y, \nu_Y)$ for ρ (see 11.23 in Royden (1968)). For given $\varepsilon > 0$, now let $g' \in L^1(Y, \nu_Y)$ be continuous and such that

$$\|g - g'\|_1 = \int |g - g'| d\nu_Y < \varepsilon$$

(using a combination of results 11.26, 14.6, and 8.5 in Royden (1968)).

Let $h: Y \times Z \rightarrow \mathbb{R}$ be an arbitrary continuous function with $|h(y, z)| < 1$ for all (y, z) . Then

$$\begin{aligned} & \left| \int h(y, z) d[\rho \otimes (\nu_n^Y - \nu^Y)](y, z) \right| \\ &= \left| \int h(y, z) g(y) d[\nu_Y \otimes (\nu_n^Y - \nu^Y)](y, z) \right| \\ &\leq \left| \int h(y, z) g'(y) d[\nu_Y \otimes (\nu_n^Y - \nu^Y)](y, z) \right| \\ &\quad + \left| \int h(y, z) [g(y) - g'(y)] d[\nu_Y \otimes (\nu_n^Y - \nu^Y)](y, z) \right|. \end{aligned}$$

The first term of this sum goes to zero because $h(y, z)g'(y)$ is continuous on $Y \times Z$ and $\nu_Y \otimes (\nu_n^Y - \nu^Y) \rightarrow 0$ weakly. The second term is bounded by ε because $\|g - g'\|_1 < \varepsilon$. Since ε is arbitrary, this means that $\int h d[\rho \otimes (\nu_n^Y - \nu^Y)] \rightarrow 0$. Since this is true for any continuous function h , we conclude that $\rho \otimes (\nu_n^Y - \nu^Y) \rightarrow 0$ weakly. This is what we wanted. Q.E.D.

For any $\alpha \in \mathcal{P}(K)$, the Riesz representation theorem implies that a probability $\int \rho d\alpha(\rho) \in \mathcal{P}(Y)$ is uniquely defined by the requirement that $\int h d[\int \rho d\alpha(\rho)] = \int [\int h d\rho] d\alpha(\rho)$ for all continuous $h: Y \rightarrow \mathbb{R}$. In words, $\int \rho d\alpha(\rho)$ is the mean of α .

LEMMA 4.2: *The function $\alpha \mapsto \int \rho d\alpha(\rho)$ from $\mathcal{P}(K)$ to $\hat{\mathcal{P}}(Y)$ is continuous (with the norm topology on K).*

PROOF: Suppose $\alpha_n \rightarrow \alpha$ in $\mathcal{P}(K)$ and let $h: Y \rightarrow \mathbb{R}$ be continuous. Since $\rho \mapsto \int h d\rho$ is weakly continuous, we have

$$\begin{aligned} \int h d \left[\int \rho d\alpha_n(\rho) \right] &= \int \left[\int h d\rho \right] d\alpha_n(\rho) \rightarrow \int \left[\int h d\rho \right] d\alpha(\rho) \\ &= \int h d \left[\int \rho d\alpha(\rho) \right]. \end{aligned}$$

Therefore $\int \rho d\alpha_n(\rho) \rightarrow \int \rho d\alpha(\rho)$ weakly. Every $\int \rho d\alpha_n(\rho)$ belongs to the convex hull of K . By a theorem of Mazur (Dunford and Schwartz (1958, p. 416)) the norm-closed convex hull of K is norm-compact. Therefore $\{\int \rho d\alpha_n(\rho)\} \subset \hat{\mathcal{P}}(Y)$ has an accumulation point in the norm, hence in the weak, topology. Since $\int \rho d\alpha(\rho)$ is the only weak accumulation point, we conclude that $\int \rho d\alpha_n(\rho) \rightarrow \int \rho d\alpha(\rho)$ in norm. Q.E.D.

PROOF: Proposition 4.2. Convex-valuedness is clear since the three conditions (i), (ii), and (iii) are linear in ν . For the closedness of the graph we only need to verify that $\Psi[\int_A P(y, \cdot) d\alpha, \nu]$ is (weakly) continuous with respect to ν, y , and α . Because of Lemma 4.1 this will follow if we establish the (norm) continuity of $(y, \alpha) \mapsto \int_A P(y, \cdot) d\alpha$.

As noted before, by the continuity hypothesis on P, K is a norm-compact subset of $\hat{\mathcal{P}}(Y)$. Denoting $P_y(\cdot) = P(y, \cdot)$, the map $(y, \alpha) \mapsto \int P(y, \cdot) d\alpha$ can be expressed as $(y, \alpha) \mapsto \int \rho d(\alpha \circ P_y^{-1})(\rho)$. Therefore, by Lemma 4.2, the desired continuity will follow from the (weak) continuity of $(y, \alpha) \mapsto (\alpha \circ P_y^{-1}) \in \mathcal{P}(K)$. To prove the latter property, let $h: K \rightarrow \mathbb{R}$ be continuous, $y_n \rightarrow y$, and $\alpha_n \rightarrow \alpha$. Then

$$\begin{aligned} &\left| \int h d(\alpha_n \circ P_{y_n}^{-1}) - \int h d(\alpha \circ P_y^{-1}) \right| \\ &= \left| \int h \circ P_{y_n} d\alpha_n - \int h \circ P_y d\alpha \right| \\ &\leq \left| \int (h \circ P_{y_n} - h \circ P_y) d\alpha_n \right| + \left| \int h \circ P_y d(\alpha_n - \alpha) \right|. \end{aligned}$$

Because $h[P_{y_n}(a)] - h[P_y(a)] \rightarrow 0$ uniformly, $\alpha_n \rightarrow \alpha$ weakly, and $a \mapsto h[P_y(a)]$ is continuous, it follows that the two terms of this sum go to zero. Hence $\alpha_n \circ P_{y_n}^{-1} \rightarrow \alpha P_y^{-1}$ weakly. Q.E.D.

PROOF: Proposition 4.4. Let $\mathcal{P}_K(Y \times Z) = \{\nu \in \mathcal{P}(Y \times Z) \mid \nu_Y \in K\}$ and note that $P^*: Y^* \times A \rightarrow \mathcal{P}(Y^*)$ can be expressed as the composition of the following

three functions:

$$\Phi_1: Y^* \times A \rightarrow Y \times A \times \mathcal{P}_K(Y \times Z)$$

defined by

$$\Phi_1(y^*, a) = (\pi(y^*), a, \Pi(y^*)),$$

$$\Phi_2: Y \times A \times \mathcal{P}_K(Y \times Z) \rightarrow K \times \mathcal{P}_K(Y \times Z)$$

defined by

$$\Phi_2(y, a, \nu) = (P(y, a), \nu),$$

$$\Phi_3: K \times \mathcal{P}_K(Y \times Z) \rightarrow \mathcal{P}_K(Y \times Z)$$

defined by

$$\Phi_3(\rho, \nu) = \Psi(\rho, \nu).$$

The function Φ_1 is measurable because Π is a transition probability, Φ_2 is continuous because of the hypothesis on P , and Φ_3 is continuous by Lemma 4.1. Therefore P^* is measurable. *Q.E.D.*

4.8. Concluding Comments

Our hypotheses are stronger than those of Rieder (1979) or, say, Nowak and Raghavan (1992). To a certain extent this is due to the topological prerequisites of our methods, but it also seems sensible to expect that the most distinctive feature of our results, namely the existence of an ergodic measure, cannot be obtained without relatively restrictive assumptions. To what extent can our hypotheses be weakened? The only significant assumptions are on the transition probabilities. It is possible that the mutual absolute continuity part could be replaced with absolute continuity of all images of P with respect to a given measure on Y . We doubt that the norm continuity can be much weakened. But we have no counterexamples.

We also lack examples where our hypotheses are satisfied and there is no Markov equilibrium (as opposed to coordinated equilibrium). A common phenomenon should, however, be kept in mind. With a continuum of states, coordination quite often appears naturally in the closure of the equilibrium (or the ε -equilibrium) set. In our model, it is also true that if $Y = [0, 1]$ (or, more generally, if Y has no isolated point), then any coordinated equilibrium payoff is the limit of ε -equilibrium payoffs as $\varepsilon \rightarrow 0$.

The stochastic games considered here have symmetric information among the players. A next logical step is the study of games with asymmetric information. Dynamic programming technique have already been used by Abreu, Pierce, Stachetti (1990) in a repeated game context with asymmetric information.

*Grad. School of Business, Stanford University, Stanford, CA 94305, U.S.A.,
Cowles Foundation, Yale University, 30 Hillhouse Ave., PO Box 2125 Yale
Station, New Haven, CT 06520, U.S.A.,*

Dept. of Economics, 326 Littauer Center, Harvard University, Cambridge, MA 02138, U.S.A.,

and

Dept. of Economics, University of Minnesota, 1035 Management and Economics, 271 19th Ave. S., Minneapolis, MN 55455, U.S.A.

Manuscript received August, 1989; final revision received December, 1993.

REFERENCES

- ABREU, D., D. PIERCE, AND E. STACHETTI (1990): "Towards a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica*, 58, 1041–1064.
- AMIR, R. (1987): "On Stochastic Games with Uncountable State and Action Spaces," unpublished, Cowles Foundation, Yale University.
- BERGIN, J., AND D. BERNHARDT (1991): "Anonymous Sequential Games with General State Space," Department of Economics, Queen's University.
- BEWLEY, T. (1972): "Existence of Equilibria in Economies with Infinitely Many Commodities," *Journal of Economic Theory*, 4, 514–540.
- BILLINGSLEY, P. (1986): *Probability and Measure*. New York: Wiley.
- BLUME, L. (1982): "New Techniques for the Study of Stochastic Equilibrium Processes," *Journal of Mathematical Economics*, 9, 61–70.
- BÖRGERS, T. (1991): "Upper Hemicontinuity of the Correspondence of Subgame-Perfect Equilibrium Outcomes," *Journal of Mathematical Economics*, 20, 89–106.
- CASS, D., AND K. SHELL (1983): "Do Sunspots Matter?" *Journal of Political Economy*, 91, 193–207.
- CHIAPPORI, P.-A., AND R. GUESNERIE (1992): "Sunspot Equilibria in Sequential Markets Models," ed. by W. Hildenbrand and H. Sonnenschein. In *Handbook of Mathematical Economics, Volume 4*. Amsterdam: North-Holland.
- DEBREU, G. (1953): "Une Economie de l'Incertain," unpublished, Electricité de France.
- DOOB, J. (1953): *Stochastic Processes*. New York: Wiley.
- DUNFORD, N., AND J. SCHWARTZ (1958): *Linear Operators, Part I: General Theory*. New York: Wiley.
- EPSTEIN, L., AND S. ZIN (1989): "Substitution, Risk Aversion, and The Temporal Behavior of Consumption and Asset Returns I: A Theoretical Framework," *Econometrica*, 57, 937–969.
- GEANAKOPOLOS, J. (1987): "Overlapping Generations," in *The New Palgrave*, ed. by J. Eatwell, M. Mieszkowski, and P. Newman. London: MacMillan.
- GOTTARDI, P. (1991): "On Stationary Monetary Equilibria in Overlapping Models with Incomplete Markets," Department of Applied Economics, Cambridge University.
- GRANDMONT, J.-M. AND W. HILDENBRAND (1974): "Stochastic Processes of Temporary Equilibrium," *Journal of Mathematical Economics*, 1, 243–277.
- GRODAL, B., AND K. VIND (1988): "Equilibrium with Arbitrary Market Structure," unpublished, Department of Economics, University of Copenhagen.
- GUESNERIE, R., AND M. WOODFORD (1992): "Endogenous Fluctuations," in *Advances in Economic Theory, Volume II*, ed. by J.-J. Laffont. Cambridge: Cambridge University Press.
- HANSEN, L. (1982): "Large Sample Properties of Generalized Methods of Moment Estimators," *Econometrica*, 50, 1029–1056.
- HANSEN, L., AND K. SINGLETON (1982): "Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models," *Econometrica*, 50, 1269–1286.
- HARRIS, C., P. RENY, AND A. ROBSON (1993): "The Existence of Subgame-Perfect Equilibrium in Continuous Games with Almost Perfect Information: a Case for Extensive-Form Correlation," Research Paper, Nuffield College, Oxford, and University of Western Ontario.
- HART, O. (1975): "On The Optimality of Equilibrium When the Market Structure is Incomplete," *Journal of Economic Theory*, 11, 418–443.
- HELLWIG, M. (1980): "Stochastic Processes of Temporary Equilibria—A Note," *Journal of Mathematical Economics*, 7, 287–299.

- HERNÁNDEZ, A., D. AND M. SANTOS (1991): "Competitive Equilibrium for Infinite-Horizon Economies with Incomplete Financial Markets," University of Wisconsin, Madison.
- (1992): "Incomplete Financial Markets with an Infinite Horizon," University of Wisconsin, Madison.
- HILDENBRAND, W. (1974): *Core and Equilibria of a Large Economy*. Princeton, New Jersey: Princeton University Press.
- HIMMELBERG, C., T. PARTHASARATHY, AND F. VAN VLECK (1976): "Optimal Plans for Dynamic Programming Plans," *Mathematics of Operations Research*, 1, 390–394.
- KELLEY, J. (1955): *General Topology*. New York: Van Nostrand.
- KIFER, Y. (1986): *Ergodic Theory of Random Transformations*. Boston: Birkhäuser.
- LEROY, S. (1973): "Risk Aversion and the Martingale Property of Stock Prices," *International Economic Review*, 14, 436–446.
- LEVINE, D., AND W. ZAME (1992): "Debt Constraints and Equilibrium in Infinite Horizon Economies with Incomplete Markets," University of California, Los Angeles, forthcoming in *Journal of Mathematical Economics*.
- LUCAS, R. (1978): "Asset Prices in an Exchange Economy," *Econometrica*, 46, 1429–1445.
- MA, C.-G. (1993): "Market Equilibrium with Heterogeneous Recursive-Utility Maximizing Agents," *Economic Theory*, 3, 243–266.
- MAGILL, M., AND M. QUINZII (1992): "Infinite Horizon Incomplete Markets," *Econometrica*, 62, 853–880.
- MAJUMDAR, M. K., AND R. SUNDARAM (1991): "Symmetric Stochastic Games of Resource Extraction: The Existence of Non-randomized Stationary Equilibrium," in *Stochastic Games and Related Topics*, ed. by T. Ferguson, T. Parthasarathy, T. E. S. Raghavan, and O. Vrieze. The Netherlands: Kluwer Academic Publishers.
- MARCEY, A., AND K. SINGLETON (1991): "Optimal Consumption and Savings Decisions and Equilibrium Asset Prices in a Model with Heterogeneous Agents Subject to Portfolio Constraints," Carnegie-Mellon University.
- MERTENS, J.-F., AND T. PARTHASARATHY (1987): "Equilibria for Discounted Stochastic Games," Research Paper 8750, CORE, Université Catholique de Louvain.
- NEVEU, J. (1965): *Mathematical Foundations of The Calculus of Probabilities*. San Francisco: Holden-Day.
- NOWAK, A. (1985): "Existence of Equilibrium Stationary Strategies in Discounted Noncooperative Stochastic Games with Uncountable State Space," *Journal of Optimization Theory and Applications*, 45, 591–603.
- (1992): "Stationary Equilibrium for Nonzero-Sum Average Payoff Ergodic Stochastic Games with General State Space," Institute of Mathematics, Technical University of Wrocław.
- NOWAK, A., AND T. E. S. RAGHAVAN (1988): "Existence of Stationary Equilibria for Discounted Stochastic Games in Arbitrary State Spaces," Statistical Lab, Technical Report No. 88-03, Univ. of Illinois at Chicago.
- (1992): "Existence of Stationary Correlated Equilibria with Symmetric Information for Discounted Stochastic Games," forthcoming in *Mathematics of Operations Research*.
- PARTHASARATHY, T. (1982): "Existence of Equilibrium Stationary Strategies in Discounted Stochastic Games," *Sankhya*, 44, 114–127.
- PARTHASARATHY, T., AND S. SINHA (1989): "Existence of Stationary Equilibrium Strategies in Non-Zero-Sum Discounted Stochastic Games with Uncountable State Space and State Independent Transitions," *International Journal of Game Theory*, 18, 189–194.
- PELED, D. (1985): "Stochastic Inflation and Government Indexed Bonds," *Journal of Monetary Economics*, 15, 123–135.
- PRESCOTT, E., AND R. MEHRA (1980): "Recursive Competitive Equilibrium: The Case of Homogeneous Households," *Econometrica*, 48, 1365–1379.
- RADNER, R. (1972): "Existence of Equilibrium of Plans, Prices and Price Expectations in a Sequence of Markets," *Econometrica*, 40, 289–303.
- REVUZ, D. (1975): *Markov Chains*. Amsterdam: North-Holland.
- RIEDER, U. (1979): "Equilibrium Plans for Nonzero Sum Markov Games," in *Game Theory and Related Topics*, ed. by O. Moeschlin and D. Pallasche. Amsterdam: North-Holland.
- ROYDEN, H. (1968): *Real Analysis (Second Edition)*. New York: Macmillan.
- SAMUELSON, P. (1958): "An Exact Consumption Loan Model of Interest with or without the Social Contrivance of Money," *Journal of Political Economy*, 66, 467–482.
- SARGENT, T. (1987): *Dynamic Macroeconomic Theory*. Cambridge, Massachusetts: Harvard University Press.

- SCHEINKMAN, J. (1980): "Note on Asset Trading in an Overlapping Generations Model," mimeo, University of Chicago.
- SHAPLEY, L. (1953): "Stochastic Games," *Proceedings of the National Academy of Sciences, U.S.A.*, 39, 1095-1100.
- SPEAR, S., AND S. SRIVASTAVA (1986): "Markov Rational Expectations Equilibrium in a Stochastic Overlapping Generations Model," *Journal of Economic Theory*, 38, 35-62.
- SPEAR, S., S. SRIVASTAVA, AND M. WOODFORD (1990): "On The Structure of Equilibrium in Stochastic Overlapping Generations Models," *Journal of Economic Theory*, 50, 265-284.
- TIROLE, J. (1985): "Asset Bubbles and Overlapping Generations," *Econometrica*, 53, 1499-1528.
- TWEEDIE, R. (1975): "Sufficient Conditions for Ergodicity and Recurrence of Markov Chains on a General State Space," *Stochastic Processes and Their Applications*, 3, 385-403.
- WANG, Y. (1990): "On the Existence of Stationary Equilibria in an OLG Model with Correlated Production Shocks," Brown University.
- WHITT, W. (1980): "Representation and Approximation of Noncooperative Sequential Games," *SIAM Journal of Control and Optimization*, 18, 33-48.