Econ. Theory 5, 461–489 (1995)

Economic Theory

Anonymous sequential games: Existence and characterization of equilibria★

J. Bergin and D. Bernhardt

Department of Economics, Queen's University, Kingston, Ontario, CANADA K7L 3N6

Received: December 16, 1992; revised version February 14, 1994

Summary. In this paper we consider Anonymous Sequential Games with Aggregate Uncertainty. We prove existence of equilibrium when there is a general state space representing aggregate uncertainty. When the economy is stationary and the underlying process governing aggregate uncertainty Markov, we provide Markov representations of the equilibria.

1 Introduction

This paper develops a framework in which dynamic games featuring both individual stochastic heterogeneity and aggregate uncertainty can be analyzed tractably. For the class of dynamic games considered – anonymous sequential games with aggregate

★ We wish to acknowledge very helpful conversations with C. d'Aspremont, B. Lipman, A. McLennan and J-F. Mertens. The financial support of the SSHRCC and the ARC at Queen's University is gratefully acknowledged. This paper was begun while the first author visited CORE. The financial support of CORE and the excellent research environment is gratefully acknowledged. The usual disclaimer applies.

Table of notation:

Λ: Agents' characteristics space ($\alpha \in \Lambda$). A: Action space of each agent ($a \in A$); Y: $Y = \Lambda \times A$; μ: Aggregate distribution on agents' characteristics; $\mathscr{M}(X)$: Space of probability measures on X; $\mathscr{C}(X)$: Space of continuous functions on X; $\mathscr{B}_X$: Family of Borel sets of X; Θ: State space of aggregate uncertainty ($\theta \in \Theta$); Θ^∞: $\Theta^\infty \equiv \times_{t=1}^{\infty} \Theta$ aggregate uncertainty for the infinite game; θ^∞: $\theta^\infty = (\theta_1, \theta_2, \ldots, \theta_t, \ldots) \in \Theta^\infty$; θ^t: $\theta^t \equiv (\theta_1, \theta_2, \ldots, \theta_t)$; $\mathbf{L}_1(\Theta^t, \mathscr{C}(\Lambda \times A), \nu_t)$: Normed space of measurable functions from Θ^t to $\mathscr{C}(\Lambda \times A)$; $\mathscr{F}(\Theta^t, \mathscr{M}(\Lambda \times A))$: Space of measurable functions from Θ^t to $\mathscr{M}(\Lambda \times A)$; X^t: $X^t = \times_{s=1}^{t} X$; $\mathscr{B}_X^t$: Borel field on X^t; ν: Distribution on Θ^∞; ν_t: Marginal distribution of ν on Θ^t; $\nu(\theta^t)(\nu(\bullet|\theta^t))$: Conditional distribution on Θ^∞ given θ^t; $\nu_t(\theta^s)(\nu_t(\bullet|\theta^s))$: Conditional distribution on Θ^t given θ^s (where $s < t$); τ_t: "Period t" distributional strategy; τ: Distributional strategy for all periods $\tau = (\tau_1, \tau_2, \ldots, \tau_t, \ldots)$; ξ_t: Transition process for agents' types; $\mathbf{P}_{t+1}(\tau_t, \theta^t, y)(\mathbf{P}_{t+1}(\bullet, \tau_t, \theta^t, y))$: Transition function associated with ξ_t; u_t: Utility function; $V_t(\alpha, a, \tau, \theta^t)$: Value function for each collection $(\alpha, a, \tau, \theta^t)$; $W_t(\alpha, \tau, \theta^t)$: Value function given optimal action a; $\mathbf{C}(\tau)$: Consistency correspondence. Distributions consistent with τ and characteristics transition functions, $\mathbf{P}_t$; $\mathbf{B}(\tau)$: Best response correspondence (which also satisfy consistency); $\mathbf{E}(\mu)$: Set of equilibrium distributional strategies; $\mathscr{M}_\infty$: $\times_{t=1}^{\infty} \mathscr{F}(\Theta^t, \mathscr{M}(\Lambda \times A))$; $\mathbf{S}$: Expanded state space for Markov construction; $v(\alpha, a, \theta)$: Value function for Markov construction; $\mathbf{P}(\tau_t^*, \theta_t, y)(\mathbf{P}(\bullet, \tau_t^*, \theta_t, y))$: Invariant characteristics transition function for Markov game.

uncertainty – we provide a set of equilibrium existence and Markov characterization results for general state spaces representing aggregate uncertainty. This class of multistage (sequential) games features a continuum of heterogeneous agents, and is characterized by the "anonymity" property that an agent's payoff in any period depends on what other agents do *only* through the aggregate distributions over agent types and their actions. These games are very natural for modeling economies where agents are "small", for example competitive economies.

The framework presents an attractive alternative to representative agent models, permitting one to address economic problems where individual stochastic heterogeneity is an important feature. In many economic environments, heterogeneity is important. Agents who differ in their abilities, endowments or preferences may make different employment decisions, hold different portfolios or purchase different goods; firms which differ in their costs may make different investment or R&D decisions, and so on. These differences can help explain both the individual allocation of resources over time, as well as the evolution of aggregate economic variables. For each agent, the dynamic evolution of such characteristics is invariably stochastic in nature: how successful was a firm's R&D investment? what was the return on an asset? what was the worker-firm match quality? what was a firm's cost shock?, etc. Questions of this sort can sometimes be addressed in the framework of the anonymous sequential game with no aggregate uncertainty (see, for instance, Jovanovic 1982, Hopenhayn 1992, Jovanovic and MacDonald 1988).

However, for many economic problems it is too restrictive to impose the requirement that the aggregate distribution of agents evolves nonstochastically. "Aggregate uncertainty" arises when the aggregate distribution on the space of agents evolves stochastically over time. For instance, modeling of business cycles demands consideration of aggregate demand shocks which affect all firms directly. Government policy choices, such as the rate of money growth, which are random from the perspective of individual agents are aggregate in nature. Technology shocks reflecting global innovations such as computers, are aggregate in nature, as are the so-called "oil" shocks. In such cases the stochastic evolution of the economy-wide aggregates is an important determinant of agent decisions and hence of economic behavior. Anonymous sequential games with aggregate uncertainty allow one to model such phenomena. For instance, Bergin and Bernhardt (1993) employ this framework to examine entry, exit, investment and R&D decisions of firms whose costs evolve stochastically and who face aggregate business cycle demand shocks.

Jovanovic and Rosenthal (1988) formally define anonymous sequential games for the case where there is no aggregate uncertainty, provide an existence theorem and illustrate its broad application to economic problems. Bergin and Bernhardt (1992) show how anonymous sequential games can be formulated with aggregate uncertainty, and provide an existence result for the case where aggregate uncertainty can be represented by a countable state space. The restriction to a countable state space is a significant shortcoming in that many problems are more easily modeled when the state space is a continuum (such as price in a market). One of the contributions of this paper is to remedy that shortcoming. The paper makes contributions along two dimensions. First, we extend the analysis in Bergin and

Bernhardt (1992) to allow for aggregate uncertainty with general state spaces. Second, we provide a set of Markov characterizations of equilibria when the model is stationary and the underlying stochastic processes in the model are Markov. In the remainder of section we briefly describe the notions of "aggregate uncertainty" and "no aggregate uncertainty" and then summarize the main results of the paper.

Let the set of agents or agent types be Λ and the action space A (common to all agents). An aggregate distribution over agent types[1] and actions is some distribution τ on $Y = \Lambda \times A$ called a distributional strategy.[2] The anonymity assumption says that the behavior of other agents affects agent α's utility only through τ. Agents' characteristics or types (e.g. technology quality) can evolve stochastically over time, so that a particular $\alpha \in \Lambda$ (the characteristics space) is not identified with "the same" player over time. At time t, if agent $\alpha \in \Lambda$ takes action $a \in A$, and the distributional strategy is τ_t, then he obtains utility $u_t(\alpha, a, \tau_t)$. Given α, τ_t and a, the player then draws a new characteristic ξ_{t+1} (reflecting idiosyncratic risk) from a distribution $\mathbf{P}_{t+1}(\tau_t, \alpha, a)$ on Λ (determining his type in period $t+1$). Sometimes, to make explicit the fact that $\mathbf{P}_{t+1}(\tau_t, \alpha, a)$ is a distribution, we will write $\mathbf{P}_{t+1}(\bullet; \tau_t, \alpha, a)$.[3] In turn, in period $t+1$ the player obtains a new characteristic, drawn from a distribution $\mathbf{P}_{t+2}(\tau_{t+1}, \xi_{t+1}, a)$, when action a is taken, the period $t+1$ distributional strategy is τ_{t+1}, and so on. Thus, idiosyncratic risk arises because a player's payoff depends on the random evolution of his characteristic in Λ space. For example, the set of agent types might be the set of possible firm technologies. The firms might have to select output and R&D actions, with the firm's technology quality evolving stochastically over time, depending on the current quality and the firm's R&D choice. The economy is competitive so that actions of competing firms affect a firm's payoffs only through the equilibrium price.

"No aggregate uncertainty" is formulated in this model as the non-stochastic evolution of a sequence of joint distributions, $\{\mu_t\}$, on the characteristics space Λ. Given the aggregate distribution, τ_t, on $\Lambda \times A$, the "no aggregate uncertainty" hypothesis means that next period's distribution over characteristics space is given by

$$\mu_{t+1}(\bullet) = \int \mathbf{P}_{t+1}(\bullet; \tau_t, \alpha, a)\tau_t(d\alpha \times da).$$

Even though each agent faces individual uncertainty through $\mathbf{P}_{t+1}$, this uncertainty at the individual level "washes out" in the aggregate so that μ_{t+1} is

[1] We identify an agent with his type or characteristic in the sense that we will refer to both agent type α and agent α interchangeably – e.g. firm α has technology type α.

[2] The properties and use of such strategies are discussed further in MasColell (1984), Jovanovic and Rosenthal (1988) and the references cited there.

[3] A "$\bullet$" is sometimes used as an argument of a measure to denote an arbitrary measurable set in the relevant space. Given two measures, μ and φ on some sigma field $\mathscr{B}$, the expression $\mu(\bullet) = \varphi(\bullet)$ means $\mu(B) = \varphi(B)$, $\forall B \in \mathscr{B}$. Given a metric space X, $\mathscr{B}_X$ is the associated sigma field.

non-stochastic. This "washing out" of individual risk is intimately related to the fact that the model has a continuum of agents (see Feldman and Gilles 1985).[4]

To illustrate the potential difficulties involved in using models with aggregate uncertainty, consider a situation where α indexes a firm's technology and where $\tilde{\xi}_\alpha$ is firm α's technology next period. One might anticipate that, from an economic perspective, agent α is better off "drawing" a good technology (a high value of $\tilde{\xi}_\alpha$). However, with a stochastic aggregate distribution and the unavoidable correlation of draws across the α's, conditional on a high value of $\tilde{\xi}_\alpha$, the distribution of technologies may be more likely to be concentrated on good technologies. Thus, in a competitive situation, given that $\tilde{\xi}_\alpha$ is high (more efficient), other firms are more likely to be more efficient and the "gain" to $\tilde{\xi}_\alpha$ of being more efficient may be offset by the fact that the competition is stiffer. Thus the expected result – that the expected payoff given greater efficiency is higher – may be reversed. It is worth noting that the difficulty here is not correlation of characteristics across agents, but of correlation between each agent's characteristic and the aggregate distribution.

Bergin and Bernhardt (1992) develop a useful decomposition of uncertainty into aggregate and idiosyncratic components. Aggregate uncertainty is introduced by having a random variable $\theta \in \Theta$ represent an aggregate "shock" to both payoffs and the transition function governing individual risk. Idiosyncratic uncertainty is represented by a second stochastic component. In the present notation, such a procedure is equivalent to writing $\eta = (\omega, \theta) \in (\Omega, \Theta)$, where θ represents aggregate uncertainty and ω embodies "idiosyncratic risk". As before, the aggregate distribution is a random variable, but if we impose the "no aggregate uncertainty" hypothesis conditional on θ, then next period's aggregate distribution is non-

[4] The "no aggregate uncertainty" hypothesis is formalized as follows: Underlying the transition function $\mathbf{P}(\bullet; \tau, \alpha, a)$ (and ignoring time subscripts) is a probability space $(N, \mathscr{B}_N, \mathbf{p})$. The process governing the evolution of individual characteristics is $\xi(\eta, \alpha, a, \tau)$: if $\eta \in N$ is drawn, agent α takes action a, and the current joint distribution on actions and agent characteristics is τ, then agent α's characteristic next period is $\xi(\eta, \alpha, a, \tau)$. The transition function is determined by this process according to:

$$\mathbf{P}(B; \tau, \alpha, a) = \mathbf{p}(\{\eta \mid \xi(\eta, \alpha, a, \tau) \in B\}), \forall B \in \mathscr{B}_\Lambda.$$

For any given η and τ, the aggregate distribution next period is given by

$$\mu^\eta(B) = \tau(\{(\alpha, a) \mid \xi(\eta, \alpha, a, \tau) \in B\}), \forall B \in \mathscr{B}_\Lambda.$$

In general, μ^η is a random measure. Letting $\mathscr{M}(\Lambda)$ denote the set of probability measures on Λ and $\mathscr{B}_{\mathscr{M}(\Lambda)}$ denote the Borel field on $\mathscr{M}(\Lambda)$, the joint distribution on $\mathscr{M}(\Lambda) \times \Lambda$ is given by:

$$\psi(Q) \equiv \mathbf{p}(\{\eta \mid (\mu^\eta, \xi(\eta, \alpha, a, \tau)) \in Q\}), Q \in \mathscr{B}_{\mathscr{M}(\Lambda)} \otimes \Lambda.$$

Similarly, the distribution of μ^η is given by:

$$\psi_\mu(Q) \equiv \mathbf{p}(\{\eta \mid \mu^\eta \in Q\}), Q \in \mathscr{B}_{\mathscr{M}(\Lambda)}.$$

The hypothesis of "no aggregate uncertainty" is the hypothesis that the distribution of this random measure, ψ_μ, is degenerate: $\exists \mu^* \in \mathscr{M}(\Lambda)$, $\mu^\eta = \mu^*$, $\mathbf{p}$ a.e. η. In this case $\mathbf{p}$ a.e. η, $\mu^*(B) = \int \mathbf{P}(B; \tau, \alpha, a)\, \tau(d\alpha \times da)$, $\forall B \in \mathscr{B}_\Lambda$. Aggregate uncertainty may be defined (by default) as the case where μ^η has a nondegenerate distribution (ψ_μ).

stochastic, conditional on θ.[5] No aggregate uncertainty conditional on θ, implies that the aggregate distribution next period, μ^*, can be computed according to:

$$\mu^*(B) = \int_Y \mathbf{P}(B; \tau, \theta, \alpha, a)\tau(d\alpha \times da), \forall B \in \mathscr{B}_\Lambda.$$

In this formulation the aggregate shock θ enters as an argument of the transition function, affecting each agent and represents the aggregate uncertainty facing every agent. Agents' actions can be conditioned on the aggregate shock, so that the aggregate shock can also affect the transition to future states indirectly through current agents' actions. Finally, θ can enter payoffs directly. For instance, θ may be an aggregate demand shock or an aggregate inflation shock which affects all firms (directly through their profits and indirectly through their actions and the future evolution of their costs). Observe that conditional on θ, the aggregate distribution and each agent's ξ realization are independent, because conditional on θ the aggregate distribution is nonstochastic. This formulation of aggregate uncertainty is discussed at some length in Bergin and Bernhardt (1992), where an existence theorem is given in the case where the state space of aggregate uncertainty is countable (i.e. Θ is countable).

Here, we use this formulation of aggregate uncertainty to provide a general equilibrium existence theorem. The extension to more general state spaces for aggregate uncertainty is important because uncountable state spaces arise in a natural way in many applications. For example, an aggregate demand shock θ shifting the intercept of a demand curve is most naturally modeled as drawn from some continuous distribution. The proof of existence of equilibrium in the general case is of independent interest because the approach used in the countable case does not carry over.[6]

We then assume that the model is stationary and provide two results on the existence of Markov equilibria. The Markov representations provide an alternative way of viewing equilibria, and the additional structure facilitates the study of equilibria, simplifying the interpretation and analysis of equilibrium behavior. The results here are closely related to some of the literature in stochastic games. In a stochastic game, a state space S is specified. There is a finite number of players, with action space $A_i(s)$, for player i, $i = 1, \ldots, n$. Let $A(s) = \times_{i=1}^n A_i(s)$. If at time t in

[5] In this formulation, the underlying probability space has the form $(\Omega \times \Theta, \mathscr{B}_\Omega \otimes \mathscr{B}_\Theta, \vartheta \otimes v) = (N, \mathscr{B}_N, \mathbf{p})$ and $\xi(\eta, \alpha, a, \tau) = \xi((\omega, \theta), \alpha, a, \tau)$. Thus, if the "aggregate shock" is θ, the aggregate distribution τ, and agent α, takes action a, then agent α's characteristic next period is drawn from the distribution $\mathbf{P}(\bullet; \tau, \theta, \alpha, a)$. The aggregate distribution is defined (for a given (ω, θ)) as:

$$\mu_{(\omega,\theta)}(B) = \tau(\{\alpha, a) | \xi((\omega, \theta), \alpha, a, \tau) \in B\}), \forall B \in \mathscr{B}_\Lambda.$$

No aggregate uncertainty conditional on θ is the requirement that, given θ, $\exists \mu^*$, such that $\mu_{(\omega,\theta)} = \mu^*$, ϑ a.e. ω. At the same time, individual agents face individual uncertainty through ω because the distribution over agent α's characteristics is given by:

$$\mathbf{P}(B; \tau, \theta, \alpha, a) = \vartheta(\{\omega | \xi((\omega, \theta), \alpha, a, \tau) \in B\}).$$

[6] The mathematical arguments developed to prove existence with the countable stage space do not extend to more general state spaces because the construction in the countable case involves selecting each finite history of aggregate shocks and developing "pointwise" arguments there.

state s_t, agents choose an action vector $a \in A(s_t)$, then the payoff to agent i is $u_i(s_t, a)$. Following the choice of $a \in A(s_t)$, a new state is drawn from some distribution $\mathbf{p}(d\tilde{s}|s_t, a)$. When agents select actions from $A(s_t)$, they observe the history of states as well as the current state, $(s_1, s_2, \ldots, s_t)$, and the history of actions $(a_1, a_2, \ldots, a_{t-1})$, where $a_\tau \in A(s_\tau)$. Payoffs are discounted over time at the rate δ, so that the present value of i's payoff at time t is $(1-\delta)\delta^{t-1}u_i(s_t, a_t)$, where $s_t \in S$ and $a_t \in A(s_t)$. Thus, a strategy for i, $\sigma_i = (\sigma_{i1}, \sigma_{i2}, \ldots, \sigma_{it}, \ldots)$, is a collection of functions with $\sigma_{it}(s_1, \ldots, s_t, a_1, \ldots, a_{t-1}) \in A_i(s_t)$. A strategy, σ_i, is called *Markov* if for all t, $\sigma_{it}(s_1, \ldots, s_t, a_1, \ldots, a_{t-1}) = \sigma^*_{it}(s_t)$, for all $(s_1, \ldots, s_t, a_1, \ldots, a_{t-1})$. If, in addition, the functions σ^*_{it} and $\sigma^*_{i\tau}$ agree, $\forall t, \tau$, then the strategy is called a *stationary Markov strategy*. In this model, when the state space S is not finite, a proof of existence of equilibrium is very difficult, and requires relatively strong assumptions on the transition probabilities on the state space. Mertens and Parthasarathy (1988) provide such a proof and also discuss some of the difficulties involved in obtaining Markov equilibrium strategies. Duffie, Geanakoplos, MasColell and MacLennan (1994) (henceforth DGMM) also discuss stochastic games as an application of a general result on existence of equilibrium. They prove existence of a stationary ergodic Markov equilibrium on an enlarged state space which includes payoffs. This circumvents some of the difficulties involved in obtaining Markov results on the S state space.

The first result we give on Markov equilibria for anonymous sequential games assumes that the stochastic process governing the θ process is Markov. In this result we enlarge the space, $\mathcal{M}(\Lambda) \times \Theta$, to include payoffs and provide a Markov characterization of equilibrium strategies. Thus, a state is a triple (μ, θ, v), where μ is a distribution over agents, θ is an aggregate shock and v a real-valued measurable function on Λ, assigning a payoff to each agent. This approach is analogous to that in DGMM who also include payoffs in the state space: the "state" at time t includes the present value of future payoffs. In a sense, this representation has a natural interpretation as a type of rational expectations equilibrium. It is worth stressing that we show that *every* equilibrium payoff in the game arises as the payoff to an equilibrium of this form. In this result, the transition functions are assumed to satisfy a form of weak* continuity whereas DGMM assume a stronger form of continuity (that the transition functions converge on Borel sets). In addition, we require no assumptions concerning absolute continuity of the transitions functions either relative to each other or relative to any fixed measure.

A possible deficiency of this Markov characterization is that the enlarged state space can make it difficult to "pin down" behavior. In order to provide a Markov result on the "natural" state space we drop the conditional no aggregate uncertainty hypothesis and return to a general model of aggregate uncertainty as a random measure μ^η. In this model, where aggregate shocks are not explicitly formulated, the "natural" state space is $\mathcal{M}(\Lambda)$. Given an underlying stationary Markov stochastic environment we demonstrate that a Markovian equilibrium exists on the standard (i.e. not enlarged) state space. This result requires stronger, but still standard, continuity assumptions on the transition functions which are similar to those in DGMM. We now turn to a description of the game. Sections 3 and 4 detail the results. Section 5 discusses the formulation of strategies. Section 6 concludes with a description of potential applications. Inconsistent with location of notation in foot 1. The appendix contains all proofs.

2 The model

The set of agents is denoted Λ with representative element α, where Λ is assumed to be a compact metric space. Λ is the "characteristics" space. Similarly, the set of actions available to any agent α is a compact metric space A. Let $Y \equiv \Lambda \times A$. An aggregate distribution on agents' characteristics is a measure μ on Λ. Given a metric space X, the set of probability measures on X is denoted $\mathscr{M}(X)$, the set of continuous functions on X is written $\mathscr{C}(X)$, and the family of Borel sets of X is given by B_X. The t-fold product of the set X is denoted $X^t = \times_{s=1}^{t} X$. We assume that the initial measure of agents is 1, i.e., $\mu(\Lambda) = 1$.

The state space representing aggregate uncertainty each period is a metric space Θ, with $\theta \in \Theta$. In the infinite period model the state space is $\Theta^\infty \equiv \times_{t=1}^{\infty} \Theta$, with representative element $\theta^\infty = (\theta_1, \theta_2, \ldots, \theta_t, \ldots)$. Let $\theta^t = (\theta_1, \theta_2, \ldots, \theta_t) \in \Theta^t = \times_{s=1}^{t} \Theta$, the history of aggregate shocks up to the end of time t. Fix an exogenously given distribution ν on Θ^∞. Denote its marginal distribution on Θ^t by ν_t, its conditional distribution on Θ^∞ given the first t elements of θ^∞ by $\nu(\theta^t)$, and the conditional distribution on Θ^t given $\theta^s, s < t$, by $\nu_t(\theta^s)$. Sometimes, for clarity of exposition we may write $\nu(\bullet \mid \theta^t)$ for $\nu(\theta^t)$, and $\nu_t(\bullet \mid \theta^s)$ for $\nu_t(\theta^s)$. The Borel field on Θ^t is denoted $\mathscr{B}^t$.

In the absence of aggregate uncertainty, a strategy is a sequence $\tau = (\tau_1, \tau_2, \ldots, \tau_t, \ldots)$, with $\tau_t \in \mathscr{M}(\Lambda \times A)$. With aggregate uncertainty, the aggregate shock history θ^t is observed when agents choose actions at time t. There are two possible ways in which the definition of a strategy might be generalized to this case: (1) define the period t strategy as a function from Θ^t to $\mathscr{M}(\Lambda \times A)$, or (2) define the period t strategy as a measure τ_t on $(\Theta^t \times \Lambda \times A)$. We adopt the first approach because there are difficulties of interpretation with the second approach and because the first approach allows a more general equilibrium existence result. These issues are discussed in greater detail in Section 5.

A period distributional strategy at time t, τ_t, is a measurable function from the space of aggregate shock histories Θ^t to $\mathscr{M}(\Lambda \times A)$, specifying for each shock history a joint distribution over agents and actions. Hence, $\tau_t \in \mathscr{F}(\Theta^t, \mathscr{M}(\Lambda \times A))$, the set of measurable functions from Θ^t to $\mathscr{M}(\Lambda \times A)$. Given τ_t and the aggregate shock history $\theta^t, \tau_t(\theta^t)$ denotes a distribution on $(\Lambda \times A)$, while $\tau(\bullet, A; \theta^t)$ denotes the corresponding marginal distribution on Λ. At time t, some aggregate shock sequence $\theta^t \in \Theta^t$ is observed[7] by agents who then choose actions: this formulation reflects the information available to agents at time t. A distributional strategy for the infinite period model is a vector $\tau = (\tau_1, \tau_2, \ldots, \tau_t, \ldots)$ of the period distributional strategies.

[7] The analysis is essentially unchanged if we allow agents to observe only the history of θ's up to the previous period so that $\tau_{t+1} \in \mathscr{F}(\Theta^t, \mathscr{M}(\Lambda \times A))$.

The evolution of agents' characteristics is described by a process $\{\xi_t\}_{t\geq 1}$. When agent α chooses the action a, the aggregate shock history is θ^t, and the period distributional strategy at this shock is $m_t = \tau_t(\theta^t)\in\mathcal{M}(\Lambda\times A)$, then α's period $t+1$ type is drawn from the distribution $\mathbf{P}_{t+1}(m_t,\theta^t,y)$, $y=(\alpha,a)$. At time t, if θ^t is the aggregate shock history and m_t is the period distributional strategy on $Y=\Lambda\times A$ at history θ^t, then the aggregate distribution at time $t+1$ is given by

$$\mu_{t+1}(\bullet)=\int_Y \mathbf{P}_{t+1}(\bullet;m_t,\theta^t,y)\tau(dy).$$

This connects the distribution on characteristics intertemporally: τ_t and τ_{t-1} are not independent. Given aggregate shock history θ^t, the marginal distribution of τ_t on Λ is a distribution on characteristics which must agree with the distribution implied by the transition process: the measure of agents in a given set in Λ at time t must equal the measure of agents entering that set from the previous period. We return to this issue of consistency below. Note that in this expression, μ_{t+1} depends on θ^t through the transition function: as θ^t varies, so also will μ_{t+1}. In period $t+1$, whatever the theta shock, θ_{t+1}, the distributional strategy τ_{t+1} at $\theta^{t+1}=(\theta^t,\theta_{t+1})$ must have a marginal distribution on Λ which agrees with μ_{t+1}, and since μ_{t+1} will generally depend nontrivially on θ^t, so also must τ_{t+1}.

Utility at time t is a function, u_t, from $\Lambda\times A\times\mathcal{M}(\Lambda\times A)\times\Theta^t$ to $\mathcal{R}$, where u_t is continuous on $\Lambda\times A$. The interpretation of the utility function is that if an agent of type α takes action a given the aggregate shock history θ^t and distributional strategy $\tau_t\in\mathcal{F}(\Theta^t,\mathcal{M}(\Lambda\times A))$, then the agent's utility is $u_t(\alpha,a,\tau_t(\theta^t),\theta^t)$. Utility at time t therefore depends on the aggregate distribution over $\Lambda\times A$, *conditional* on θ^t. For ease of notation we write $u_t(\alpha,a,\tau_t,\theta^t)$ with the interpretation that given $\tau_t\in\mathcal{F}(\Theta^t,\mathcal{M}(\Lambda\times A))$, u_t depends only on the value of τ_t at θ^t: $u_t(\alpha,a,\tau_t,\theta^t)\stackrel{def}{=}u_t(\alpha,a,\tau_t(\theta^t),\theta^t)$. We assume that $\forall t$, $|u_t|\leq K'<\infty$, so that without loss of generality we may take $0\leq u_t(\alpha,a,\tau_t,\theta^t)\leq K<\infty$, $\forall(t,\alpha,a,\tau_t,\theta^t)$. Payoffs are discounted: the discount rate at time t is δ_t, where $sup_{t\geq 1}\delta_t<1$, so that the present value of time t payoffs is $(\prod_{s=1}^t\delta_s)u_t$. In the stationary model, we set $\delta_t=(1-\delta)\delta^{t-1}$, where $0<\delta<1$.

The sequence of events at time t is the following. First, the period t aggregate shock, θ_t, is realized. Then agent α picks an action $a\in A$. Given the history of aggregate shocks (including the current shock) θ^t, and $\tau(\theta^t)$, agent α receives utility $u_t(\alpha,a,\tau_t(\theta^t),\theta_t)$. Following this the agent's characteristics for period $t+1$, ξ_{t+1}, are drawn from the transition distribution on Λ, $\mathbf{P}_{t+1}(\tau_t(\theta^t),\theta^t,y)$, $y=(\alpha,a)$.

2.1 Continuity assumptions on payoffs and transition functions

In this section we give the main technical assumptions. Given a Banach space $(X,\|\ \|_X)$, let $\mathbf{L}_1(\Theta^t,X,\nu_t)$ represent the set of integrable functions[8] from Θ^t to X

[8] A function f: $\Theta^t\to X$ is called simple if $\exists x_1,x_2,\ldots,x_n\in X$, $f=\sum_{i=1}^n x_i\chi E_i$, where $E_i\in\mathcal{B}^t$ and χE_i is the indicator function of E_i. The function f: $\Theta^t\to X$ is called ν_t-measurable if $\exists$ a sequence of simple functions $\{f_n\}$, such that $lim_n\|f_n-f\|=0$, ν_t-almost everywhere. Finally, a ν_t-measurable function f is called (Bochner) integrable if $\exists$ a sequence of simple functions $\{f_n\}$, such that $lim_n\int_{\Theta^t}\|f_n-f\|\nu_t(d\theta^t)=0$. See Diestel and Uhl (1977) for further details.

with norm $\int_{\Theta^t} \| f(\theta^t) - g(\theta^t) \| v_t(d\theta^t)$. For $f, g \in \mathscr{C}(Y)$, let $\| f - g \|_{\mathscr{C}(Y)} \equiv sup_y | f(y) - g(y) |$. If $f, g \in \mathbf{L}_1(\Theta^t, \mathscr{C}(Y), v_t)$, then $f(y, \theta^t)$ and $g(y, \theta^t)$ are functions on Y for each θ^t, and given θ^t, $\| f - g \|_{\mathscr{C}(Y)} \equiv sup_y | f(y, \theta^t) - g(y, \theta^t) |$ (or $\| f - g \|_{\mathscr{C}(Y)}(\theta^t)$ to make the dependence on θ^t explicit). The $\mathbf{L}_1(\Theta^t, \mathscr{C}(Y), v_t)$ norm topology is determined by the metric $\int_{\Theta^t} \| f - g \|_{\mathscr{C}(Y)} v_t(d\theta^t)$.

Define a topology on $\mathscr{F}(\Theta^t, \mathscr{M}(\Lambda \times A))$ according to the following convergence criterion:[9] Say that $\tau_t^n \to \tau_t$ if and only if for all $f \in \mathscr{C}(Y)$ and $g \in \mathbf{L}_1(\Theta^t, \mathscr{R}, v_t)$,

$$\int f(y) g(\theta^t) \tau_t^n(dy; \theta^t) v_t(d\theta^t) \to \int f(y) g(\theta^t) \tau_t(dy; \theta^t) v_t(d\theta^t).$$

This is the generalization of the weak* topology to the case of random distributions: it is the coarsest topology on $\mathscr{F}(\Theta^t, \mathscr{M}(\Lambda \times A))$ for which $\int f(y) g(\theta^t) \tau_t(dy; \theta^t) v_t(d\theta^t)$ is continuous in τ_t. With this topology, $\mathscr{F}(\Theta^t, \mathscr{M}(\Lambda \times A))$ is compact (Mertens 1986).

Given a continuous function f on Λ ($f \in \mathscr{C}(\Lambda)$), let $\int_\Lambda f(\xi) \mathbf{P}_{t+1}(d\xi, \tau_t(\theta^t), \theta^t, y)$ be denoted $\mathbf{P}_{t+1}(f, \tau_t(\theta^t), \theta^t, y)$. The transition distribution $\mathbf{P}_{t+1}(\tau_t(\theta^t), \theta^t, y)$ is assumed *weak** continuous in y: for each $f \in \mathscr{C}(\Lambda)$, $\mathbf{P}_{t+1}(f, \tau_t(\theta^t), \theta^t, y)$ is a continuous function of y. In addition, the following continuity conditions relative to τ are imposed. For fixed f, and $\tau_t \in \mathscr{F}(\Theta^t, \mathscr{M}(\Lambda \times A))$, both $\mathbf{P}_{t+1}$ and u_t may be viewed as continuous real-valued functions on Y for each θ^t and hence as elements of $\mathbf{L}_1(\Theta^t, \mathscr{C}(Y), v_t)$. To simplify notation, write $\mathbf{P}_{t+1}(f, \tau_t, \theta^t, y)$ for $\mathbf{P}_{t+1}(f, \tau_t(\theta^t), \theta^t, y)$, where τ_t is understood to be an element of $\mathscr{F}(\Theta^t, \mathscr{M}(\Lambda \times A))$. We assume that for all $f \in \mathscr{C}(\Lambda)$, $\mathbf{P}_{t+1}(f, \tau_t, \theta^t, y)$ and $u_t(\alpha, a, \tau_t, \theta^t)$ are norm continuous in τ_t:

$$\int_{\Theta^t} sup_y | \mathbf{P}_{t+1}(f, \tau_t^n, \theta^t, y) - \mathbf{P}_{t+1}(f, \tau_t, \theta^t, y) | v_t(d\theta^t) \xrightarrow{\tau_t^n \to \tau_t} 0$$

and

$$\int_{\Theta^t} sup_y | u_t(y, \tau_t^n, \theta^t) - u_t(y, \tau_t, \theta^t) | v_t(d\theta^t) \xrightarrow{\tau_t^n \to \tau_t} 0.$$

These conditions are the natural generalization of the conditions on preferences and transition probabilities in the no aggregate uncertainty case. Consider the condition on utility and suppose first that Θ contains just one element, θ^*, so that, in effect, there is no aggregate uncertainty. In this case, θ^t has only one possible value, $(\theta^*)^t$, and so can be dropped from the notation $\tau(\theta^t)$. Thus, the strategy at time t, τ_t, is simply an element of $\mathscr{M}(\Lambda \times A)$ (with the *weak** topology) and utility is a function of (α, a, τ_t). The condition on utility is then:

$$sup_y | u_t(y, \tau_t^n) - u_t(y, \tau_t) | \xrightarrow{\tau_t^n \to \tau_t} 0.$$

A sufficient condition for this to hold is that u_t be continuous on the compact space $\Lambda \times A \times \mathscr{M}(\Lambda \times A)$. Next, consider the case where Θ is countable, so that Θ^t is countable for each t. Suppose also, that $v_t(\theta^t) > 0, \forall \theta^t \in \Theta^t$. In this case, a sequence τ_t^n converges to τ_t if and only if $\tau_t^n(\theta^t)$ converges *weak**, for each $\theta^t \in \Theta^t$. The condition

[9] Thanks are due to J-F. Mertens for suggesting this topology.

on utility is then:

$$sup_y|u_t(y, \tau_t^n(\theta^t), \theta^t) - u_t(y, \tau_t(\theta^t), \theta^t)| \xrightarrow{\tau_t^n \to \tau_t} 0, \quad \forall \theta^t \in \Theta^t.$$

Given $\theta^t, u_t(y, \tau_t(\theta^t); \theta^t)$ maps $Y \times \mathscr{M}(A \times A)$ to $\mathscr{R}$. If, for each θ^t, u_t is continuous on $Y \times \mathscr{M}(A \times A)$, then, as before, the condition on utility is satisfied. A similar discussion applies to the condition on the transition probability. In the case where Θ is countable, the assumption implies that for each $\theta^t, \mathbf{P}_{t+1}(\tau_t, \theta^t, y)$ is *weak** continuous in τ_t. Thus, these assumptions are the natural generalization of *weak** continuity of utilities and transition probabilities on aggregate distributions to the case of random distributions.

For many applications the natural formulation of the transition process will not include the aggregate distribution and in such cases the norm continuity assumption on the transition probability is trivially satisfied. For example, in the model of R&D discussed in the introduction, the success of a firm's research efforts in terms of improving its technology would not be expected to depend on the R&D efforts of other firms although such efforts would affect the firm's competitive position in subsequent periods. In that case the aggregate distribution does not appear as an argument of the transition function and norm continuity is automatically satisfied.

3 Equilibrium

In this section the first result we give attaches a value function to each collection of $(\alpha, a, \tau, \theta^t)$. This value function, $V_t(\alpha, a, \tau, \theta^t)$, details the payoff to agent α at time t, when α takes action a, the aggregate shock history is θ^t and the distributional strategy is τ. The proof of the existence of a value function does not require that the time t utility function, u_t, be continuous in the aggregate shock θ_t. We then begin to set out the conditions for an equilibrium. The definition of equilibrium is somewhat involved since the set of agents "available" to optimize at any point in time must be consistent with the set of agents carried forward from the previous period through the transition function and the distribution over characteristics at that (previous) period. Furthermore, the "set of agents" is defined essentially by the distribution over characteristics, and this is a random variable. We formulate these intertemporal consistency conditions on strategies and define best response mappings. We then prove that there exists a distributional strategy consistent with itself, for which almost all agents are optimizing at almost all aggregate shock histories, so that an equilibrium exists.

In defining value functions we first consider a truncated n-period version of the game and families of value functions $V_t^n(\alpha, a, \tau, \theta^t)_{t=1}^n$ and $W_t^n(\alpha, \tau, \theta^t)_{t=1}^n$, where, for example, $W_t^n(\alpha, \tau, \theta^t)$ is the expected payoff in the truncated game (from period t on) to player α given history θ^t and τ, when α plays optimally from period t to the end of the n period game. We show that for each t, these functions are continuous in (α, a) and α respectively, and norm continuous in τ. We then demonstrate that these functions converge uniformly as $n \to \infty$, so that the limiting value functions also have these properties. All proofs are given in the appendix.

Theorem 1 *For each t, there exist value functions* $V_t(\alpha, a, \tau, \theta^t)$ *and* $W_t(\alpha, \tau, \theta^t)$, *which* *are continuous in* (α, a) *and* α *respectively, norm continuous in* τ *and satisfy* $W_t(\alpha, \tau, \theta^t) = max_\alpha V_t(\alpha, a, \tau, \theta^t)$.

These valuation functions are used below to define the "best response" correspondence.

We now formulate the appropriate *consistency conditions* on the distributional strategy sequences $\tau = \{\tau_t\}$ (in terms of the distribution over characteristics). If τ is an equilibrium distributional strategy, then the measure of agents in existence at the beginning of period t, as given by the period t distributional strategy τ_t, must coincide with the measure mapped from period $t-1$: in any equilibrium, a strategy must be consistent with itself. Note that given a distributional strategy τ, and aggregate shock history θ^t, the distribution over characteristics at time t, "implied" by τ is given by $\int_Y \mathbf{P}_t(B, \tau_{t-1}, \theta^{t-1}, y)\tau_{t-1}(dy; \theta^{t-1})$, for all Borel sets $B \in \mathscr{B}_\Lambda$. Any distribution on $\Lambda \times A$ whose marginal distribution agrees with this distribution is consistent with τ at time t. Thus, given τ, t and θ^t, there is a set of distributions $\tilde{\tau}_t(\theta^t)$ on $\Lambda \times A$, such that the marginal of $\tilde{\tau}_t(\theta^t)$ on Λ agrees with that implied by τ. The collection of such distributions (as t and θ^t vary) is the set of distributions consistent with τ.

Definition 1 *Let* $\tilde{\tau} = \{\tilde{\tau}_t\}_{t=1}^{\infty}$ *and* $\tau = \{\tau_t\}_{t=1}^{\infty}$, *with* $\tau_1(\bullet, A) = \mu_1(\bullet)$. *Say that* $\tilde{\tau}$ *is consistent with* τ *if:*

$$\int_\Theta \tilde{\tau}_1(f, A; \theta)g(\theta)\nu_1(d\theta) = \int_\Theta \mu_1(f)g(\theta)\nu_1(d\theta), \forall f \in \mathscr{C}(\Lambda), g \in \mathbf{L}_1(\Theta, \mathscr{R}, \nu_1),$$

$$\int_{\Theta^2} \tilde{\tau}_2(f, A; \theta^2)g(\theta^2)\nu_2(d\theta^2) = \int_{\Theta^2}\int_Y \mathbf{P}_{\xi_2}(f, \tau_1, \theta^1, y)\tau_1(dy; \theta^1)g(\theta^2)\nu_2(d\theta^2),$$

$$\forall f \in \mathscr{C}(\Lambda), g \in \mathbf{L}_1(\Theta^2, \mathscr{R}, \nu_2),$$

and for period t,

$$\int_{\Theta^t} \tilde{\tau}_t(f, A; \theta^t)g(\theta^t)\nu_t(d\theta^t) = \int_{\Theta^t}\int_Y \mathbf{P}_t(f, \tau_{t-1}, \theta^{t-1}, y)\tau_{t-1}(dy; \theta^{t-1})g(\theta^t)\nu_t(d\theta^t),$$

$$\forall f \in \mathscr{C}(\Lambda), g \in \mathbf{L}_1(\Theta^t, \mathscr{R}, \nu_t).$$

These conditions imply that $\tilde{\tau}_t(f, A; \theta^t) = \int_Y \mathbf{P}_t(f, \tau_{t-1}, \theta^{t-1}, y)\tau_{t-1}(dy; \theta^{t-1})$, almost everywhere θ^t (relative to ν_t). Recall, if $\mu \in \mathscr{M}(\Lambda)$ and f is a measurable function on Λ, $\mu(f)$ denotes $\int f d\mu$. Thus, the condition imposed is that the distribution over characteristics space Λ, determined by the distributional strategy $\tilde{\tau}$ at time t, $\tilde{\tau}_t(\bullet, A; \theta^t)$, is consistent, given the distributional strategy τ, with the characteristics distribution implied by the characteristics transition function, $\mathbf{P}_t(\tau_{t-1}, \theta^{t-1}, y)$, and the distribution over previous state variables determined by τ, $\tau_{t-1}(\bullet; \theta^{t-1})$.

Denote the collection of strategies which are consistent with τ by $\mathbf{C}(\tau)$. In view of definition 1, $\mathbf{C}(\tau)$ may be defined as $\mathbf{C}(\tau) \equiv \times_{t=1}^{\infty} \mathbf{C}_t(\tau)$, where

$$\mathbf{C}_1(\tau) = \{\tilde{\tau}_1 \mid \int_\Theta \tilde{\tau}_1(f, A; \theta)g(\theta)\nu_1(d\theta) = \int_\Theta \mu_1(f)g(\theta)\nu_1(d\theta), \forall f \in \mathscr{C}(\Lambda), g \in \mathbf{L}_1(\Theta, \mathscr{R}, \nu_1)\},$$

and for $t \geq 2$,

$$\mathbf{C}_t(\tau) = \left\{ \tilde{\tau}_t \middle| \int_{\Theta^t} \tilde{\tau}_t(f, A; \theta^t) g(\theta^t) v_t(d\theta^t) \right.$$
$$\left. = \int_{\Theta^t} \int_Y \mathbf{P}_t(f, \tau_{t-1}, \theta^{t-1}, y) \tau_{t-1}(dy; \theta^{t-1}) g(\theta^t) v_t(d\theta^t), \forall f \in \mathscr{C}(A), g \in \mathbf{L}_1(\Theta^t, \mathscr{R}, v_t) \right\}.$$

Norm continuity of $\mathbf{P}_t(f, \tau_{t-1}, \theta^{t-1}, y)$ in τ (viewing $\mathbf{P}_t(f, \tau_{t-1}, \theta^{t-1}, y)$ as an element of $\mathbf{L}_1(\Theta^{t-1}, \mathscr{C}(Y), v_{t-1})$, for fixed f) ensures that these equalities are continuous in τ sequences. A strategy consistent with itself is a fixed point of $\mathbf{C}$. Continuity of the equalities in τ sequences gives the following result.

Theorem 2 *The correspondence $\mathbf{C}(\tau) \equiv \times_{t=1}^{\infty} \mathbf{C}_t(\tau)$ is non-empty, upper-hemicontinuous and convex-valued.*

We now consider those distributional strategy sequences in which almost all agents are maximizing for almost all aggregate shock histories, θ^t. Consider the period t valuation function $V_t(\alpha, a, \tau, \theta^t)$. This gives the payoff to agent α if the distributional strategy is given by τ, the aggregate shock history to period t is θ^t, and α chooses a. Given τ, $\mathbf{C}_t(\tau)$ gives the set of period t distributional strategies whose marginal distributions on characteristics space agrees with the distribution over characteristics space implied by τ and the transition functions. For a strategy to be an equilibrium, we require that it be consistent (with itself) and that at every time period, at almost all histories (θ shocks), almost all agents are optimizing. If, for the moment, we fix a "representative" θ^t, then τ and the transition functions imply some distribution, say $\lambda_t(\bullet; \theta^t)$ on Λ. With agents selecting actions optimally, the payoff to α is $max_a V_t(\alpha, a, \tau, \theta^t) = W_t(\alpha, \tau, \theta^t)$. Let $h(\alpha, \theta^t)$, be an optimal choice for α (at θ^t), with h a measurable function on $\Lambda \times \Theta^t$. Then h and λ_t determine a joint distribution $\hat{\tau}_t$ on $\Lambda \times A$, for each θ^t: $\hat{\tau}_t(\theta^t)$ (or $\hat{\tau}_t(\bullet, \bullet; \theta^t)$). Note that $\hat{\tau}_t \in \mathbf{C}_t(\tau)$. By construction, if $\tilde{\tau}_t \in \mathbf{C}_t(\tau)$, for almost all θ^t, $\int_Y V_t(\alpha, a, \tau, \theta^t) \tilde{\tau}_t(dy; \theta^t) \leq \int_Y V_t(\alpha, a, \tau, \theta^t) \hat{\tau}_t(dy; \theta^t)$. This inequality follows directly from the fact that $\int_Y V_t(\alpha, a, \tau, \theta^t) \hat{\tau}_t(dy; \theta^t) = \int_Y V_t(\alpha, h(\alpha, \theta^t), \tau, \theta^t) \lambda_t(d\alpha; \theta^t)$. Further, if for all $\tilde{\tau}_t \in \mathbf{C}(\tau)$, $\int_Y V_t(\alpha, a, \tau, \theta^t) \tilde{\tau}_t(dy; \theta^t) \leq \int_Y V_t(\alpha, a, \tau, \theta^t) \hat{\tau}_t(dy; \theta^t)$, for almost all θ^t, then almost all agents are optimizing at almost all θ^t. If τ_t and $\hat{\tau}_t$ coincide in this definition, for each t and θ^t, then under the distributional strategy τ, every period almost all agents are optimizing at almost all aggregate shocks. At a representative θ^t (in the measure 1 set) $\int_Y V_t(\alpha, a, \tau, \theta^t) \hat{\tau}_t(dy; \theta^t) = \int_Y \max_{a \in A} V_t(\alpha, a, \tau, \theta^t) \hat{\tau}_t(dy; \theta^t)$. In this case, since τ is consistent with itself, it is an equilibrium.

Definition 2 *Let τ be a distributional strategy consistent with itself. Then τ is an equilibrium if for each t,*

$$sup_{\tilde{\tau} \in C_t(\tau)} \int_{\Theta^t} \int_Y V_t(\alpha, a, \tau, \theta^t) \tilde{\tau}_t(dy; \theta^t) v_t(d\theta^t) \leq \int_{\Theta^t} \int_Y V_t(\alpha, a, \tau, \theta^t) \tau_t(dy; \theta^t) v_t(d\theta^t).$$

Define a best response mapping, $\mathbf{B}(\tau)$:

$$\mathbf{B}(\tau) = \left\{ \hat{\tau} \in \mathbf{C}(\tau) | \forall t, sup_{\tilde{\tau} \in C_t(\tau)} \int_{\Theta^t} \int_Y V_t(\alpha, a, \tau, \theta^t) \tilde{\tau}_t(dy; \theta^t) v_t(d\theta^t) \right.$$

$$\leq \int_{\Theta^t} \int_Y V_t(\alpha, a, \tau, \theta^t)\hat{\tau}_t(dy; \theta^t)v_t(d\theta^t) \bigg\}.$$

A fixed point of **B** is an equilibrium. By construction, if $\tilde{\tau}$ in $\mathbf{B}(\tau)$, then for all t, for almost all θ^t, almost all agents are maximizing at state θ^t. The next theorem shows that **B** satisfies the conditions of the Glicksbuerg Fan Theorem. That is, **B** is convex-valued, non-empty and upper-hemicontinuous.

Theorem 3 *The correspondence* **B** *satisfies the conditions of the Glicksberg Fan theorem and hence has a fixed point, which is an equilibrium of the game.*

4 Markov equilibria

We now show that when the model is stationary and the θ process Markov, there exists a *Markov equilibrium.* More precisely, we show constructively that for every equilibrium, there is an (expected payoff) equivalent Markov equilibrium. This result uses the conditional no aggregate uncertainty formulation involving the θ process. The Markov equilibrium is on an enlarged state space which includes payoffs. This is similar to DGMM who also use an enlarged state space which includes payoffs. However, we impose relatively weak assumptions on the transition function of the process. We conclude this section by dropping the conditional no aggregate uncertainty hypothesis and returning to a general model of aggregate uncertainty as a random measure μ^η. In that environment, under strong continuity assumptions similar to those in DGMM, we provide a result on the existence of Markov equilibrium where the state space is just the aggregate distribution over characteristics. This illustrates an alternative approach to modeling aggregate uncertainty, and does not require that expectations enter the state space.

For the model with conditional no aggregate uncertainty we first impose the following stationarity assumptions (with a slight abuse of notation):

1. $u_t(\alpha, a, \tau_t, \theta^t) = u(\alpha, a, \tau_t, \theta_t)$: utility is time independent and depends only on the current value of θ.
2. $\mathbf{P}_t(\bullet; \tau_{t-1}, \theta^{t-1}, y) = \mathbf{P}(\bullet; \tau_{t-1}, \theta_{t-1}, y)$: the transition function is Markov.
3. $v(\theta^t)$ depends only on θ_t. With a mild abuse of notation $v(\bullet \mid \theta^t) = v(\bullet \mid \theta_t)$: the aggregate shock process is Markov.

In addition, we assume that

1. $u(\alpha, a, \tau_t, \theta_t)$ is continuous in (α, a, θ_t) and norm continuous in τ_t.
2. $\mathbf{P}(\alpha, a, \tau_t, \theta_t)$ is *weak** continuous in (α, a, θ_t) and norm continuous in τ_t.
3. $v(\bullet \mid \theta_t)$ is *weak** continuous on Θ.
4. Θ is a compact metric space.

These additional assumptions imply that the value functions $V_t(\alpha, a, \tau, \theta^t)$ and $W_t(\alpha, \tau, \theta^t)$ (given in theorem 1) are continuous in θ^t. We now introduce a state space, **S**, and define equilibrium Markov strategies relative to this state space.

Given an initial distribution μ over the characteristics space and an initial aggregate shock θ, we denote the associated set of equilibrium distributional strategies as:

$\mathbf{E}(\mu, \theta) = \{\tau \in \mathscr{M}_\infty \,|\, \tau$ *is an equilibrium of the game with initial characteristics distribution* μ *and initial aggregate shock* $\theta\}$, where $\mathscr{M}_\infty = \times_{t=0}^{\infty} \mathscr{F}(\Theta^t, \mathscr{M}(\Lambda \times A))$ and $\mathscr{F}(\Theta^0, \mathscr{M}(\Lambda \times A)) = \mathscr{M}(\Lambda \times A))$. Define the state space $\mathbf{S}$:

$$\mathbf{S} = \{(\mu, v, \theta) \in \mathscr{M}(\Lambda) \times \mathscr{C}(\Lambda) \times \Theta \,|\, \exists \tau \in \mathbf{E}(\mu, \theta) \text{ and } v(\alpha) = W_1(\alpha, \tau, \theta), \forall \alpha \in \Lambda\}.$$

Thus, $(\mu, v, \theta) \in \mathbf{S}$, means that given initial conditions (μ, θ) there is an equilibrium strategy τ, such that the expected payoff to agent α in this equilibrium is $v(\alpha)$. In addition, define a correspondence φ: $\mathbf{S} \to \mathscr{M}_\infty$ according to:

$$\varphi(\mu, v, \theta) = \{\tau \in \mathscr{M}_\infty \,|\, \tau \in \mathbf{E}(\mu, \theta), v(\alpha) = W_1(\alpha, \tau, \theta), \forall \alpha \in \Lambda\}.$$

The correspondence φ associates to any point $(\mu, v, \theta) \in \mathbf{S}$ an equilibrium strategy τ (in the game with initial characteristics distribution μ and initial aggregate shock θ), with the property that the payoff to α is $v(\alpha)$. Under the additional assumption of continuity in θ, the correspondence φ is an upper-hemicontinuous correspondence. Define a Markov equilibrium:

Definition 3 *An equilibrium distributional strategy* $\bar{\tau}$ *is a Markov Equilibrium if for almost all* θ^t, $\theta^{t'}$ *such that*

$$(i)\ \mu(\bullet \,|\, \theta^{t-1}) = \mu(\bullet \,|\, \theta^{t'-1}), \quad (ii)\ W_t(\alpha, \bar{\tau}, \theta^t) = W_{t'}(\alpha, \bar{\tau}, \theta^{t'}) \quad \text{and } (iii)\ \theta_t = \theta_{t'},$$

the strategy $\bar{\tau}$ *satisfies* $\bar{\tau}(\bullet, \bullet \,|\, \theta^t) = \bar{\tau}(\bullet, \bullet \,|\, \theta^{t'})$.

Thus, an equilibrium distributional strategy is a Markov equilibrium if the behavior at two different histories is the "same", when the distributions on characteristics, the expected payoffs to all agents, and the aggregate shocks agree. The following theorem asserts that every equilibrium payoff in the game arises as the payoff to some Markov equilibrium.

Theorem 4 *Given an equilibrium* τ *of the game with initial characteristics distribution* μ *and initial state* θ, *there is a Markov equilibrium,* $\bar{\tau}$, *such that the first period payoff to each agent is unchanged: the expected payoff to* α *is the same under* $\bar{\tau}$ *as* τ.

The proof involves taking a pointwise measurable selection, $\tau^*, \tau^*(\mu, v, \theta) \in \varphi(\mu, v, \theta)$, for all $(\mu, v, \theta) \in \mathbf{S}$ which we use to construct the Markov equilibrium $\bar{\tau}$. Future payoffs are supported by reapplying the first component of $\tau^*(\mu, v, \theta), \tau_1^*(\mu, v, \theta)$, in succeeding periods $2, 3, \ldots,$ thus introducing Markov stationarity. This follows an approach given in Bergin (1989).

For the final result, we return to a basic formulation of aggregate uncertainty. In the model with an aggregate uncertainty parameter (θ) identified explicitly, aggregate uncertainty is modeled with the aggregate distribution conditionally nonstochastic, given the current aggregate shock. Typically, however, "Markov – type" results require a degree of continuity in the transition process governing the state variable (in the sense of absolute continuity relative to a fixed measure or relative to the transition measure at all states, say). Conditional no aggregate uncertainty runs counter to this type of assumption. Reverting from the conditional no aggregate uncertainty assumption to the general specification permitting aggregate uncertainty allows us to address the issue of Markov structure with standard (although strong) assumptions on the transition process.

We return to the process $\xi(\eta, \alpha, a, \tau)$ governing the evolution of individual characteristics, modified slightly to include the distribution on characteristics, μ: $\xi(\eta, \alpha, a, \tau, \mu)$. In this case, for a given (η, τ, μ), next period's aggregate distribution is given by

$$\mu^\eta(B) = \tau(\{(\alpha, a) | \xi(\eta, \alpha, a, \tau, \mu) \in B\}), \forall B \in \mathscr{B}_A.$$

The corresponding distributions on the space of measures over characteristics are:

$$\psi_\mu(Q) \equiv \mathbf{p}(\{\eta | \mu^\eta \in Q\}), Q \in \mathscr{B}_{\mathscr{M}(A)},$$

and

$$\psi(Q) \equiv p(\{\eta | (\mu^\eta, \xi(\eta, \alpha, a, \tau)) \in Q\}), Q \in \mathscr{B}_{\mathscr{M}(A)} \otimes A.$$

Since θ is no longer separate from η, it no longer enters the utility function explicitly: utility is given by a (time independent) function, $u(\alpha, a, \tau, \mu)$. The transition function now is a distribution on $\mathscr{M}(A) \times A$, where $\mathbf{P}_{\mathscr{M}\xi}(\bullet, \bullet, \alpha, a, \tau, \mu)$ gives a distribution over $\mathscr{M}(A) \times A$, given (α, a, τ, μ). A period t distributional strategy is a function from $\mathscr{M}(A)$ to $\mathscr{M}(A \times A)$. For a fixed measure ψ on $\mathscr{M}(A)$, the natural topology on the space of these functions is given by the following criterion of convergence: $\tau^k \to \tau$ if $\forall f \in \mathscr{C}(Y), g \in \mathbf{L}_1(\mathscr{M}(A), \mathscr{M}(A \times A), \psi)$,

$$\int_Y f(y) \tau^k(dy; \mu) g(\mu) \psi(d\mu) \to \int_Y f(y) \tau(dy; \mu) g(\mu) \psi(d\mu).$$

We make the following assumptions. There is a fixed measure ψ on $\mathscr{M}(A)$ such that

1. u is continuous in (α, a) and norm continuous in τ (relative to ψ):

$$\int_{\mathscr{M}(A)} sup_y | u(y, \tau^k, \tilde{\mu}) - u(y, \tau^k, \tilde{\mu}) | \psi(d\tilde{\mu}) \xrightarrow{\tau_t^k \to \tau_t} 0.$$

2. $\mathbf{P}_{\mathscr{M}\xi}$ is continuous in (α, a) and norm continuous in τ (relative to ψ) on *measurable functions*, f, on $\mathscr{M}(A) \times A$:

$$\int_{\mathscr{M}(A)} sup_y | \mathbf{P}_{\mathscr{M}\xi}(f, y, \tau^k, \tilde{\mu}) - \mathbf{P}_{\mathscr{M}\xi}(f, y, \tau, \tilde{\mu}) | \psi(d\tilde{\mu}) \xrightarrow{\tau_t^k \to \tau_t} 0.$$

3. The μ component of the transition functions is dominated by ψ uniformly, $\exists b < \infty$ such that for any $(y, \tau, \tilde{\mu})$, $\exists f$ measurable, f: $\mathscr{M}(A) \to \mathscr{R}, 0 \le f \le b$, such that

$$\mathbf{P}_{\mathscr{M}\xi}(X, A, y, \tau, \tilde{\mu}) = \int_X f(\mu) \psi(d\mu).$$

In this model, the state space for the Markov formulation is $\mathscr{M}(A)$. An equilibrium is Markov if the current distributional strategy τ_t depends only on the current state, μ_t. The intertemporal consistency conditions on the distributions have the form: $\forall f \in \mathscr{C}(\mathscr{M}(A)), \forall g \in \mathbf{L}_1(\mathscr{M}(A), \mathscr{R}, \psi)$

$$\int_{\mathscr{M}(A)} \hat{\tau}_{t+1}(f, A; \mu) g(\mu) \psi(d\mu) = \int \mathbf{P}_{\mathscr{M}\xi}(f, A; y, \tau_t, \mu) g(\mu) \psi(d\mu), \forall y \in Y, t \ge 1,$$

and,

$$\int_{\mathcal{M}(A)} \hat{\tau}_1(f, A; \mu) g(\mu) \psi(d\mu) = \int_{\mathcal{M}} \mu(f) g(\mu) \psi(d\mu).$$

In this case, say that $\hat{\tau}$ is consistent with τ.

The proof that a Markov equilibrium exists proceeds along similar lines to theorems 1 and 3. As before, this entails establishing the existence of valuation functions $\{V_t(\alpha, a, \tau, \mu)\}_{t \geq 1}$ and $\{W_t(\alpha, \tau, \mu)\}_{t \geq 1}$, where $\tau = (\tau_1, \tau_2, \ldots,)$, $\tau_t: \mathcal{M}(A) \to \mathcal{M}(A \times A)$. We then define intertemporal consistency conditions for the distributional strategies and show that the consistency and best response mappings satisfy the conditions of the Glicksberg Fan Theorem.

As before, we prove the existence of value functions by looking at a truncated n-period version of the game, establishing the existence of $\{V_t^n(\alpha, a, \tau, \mu)\}_{t \geq 1}$ and $\{W_t^n(\alpha, \tau, \mu)\}_{t \geq 1}$, and then take limits to obtain $\{V_t(\alpha, a, \tau, \mu)\}_{t \geq 1}$ and $\{W_t(\alpha, \tau, \mu)\}_{t \geq 1}$. Note that even if the function $W_t(\alpha, \tau, \mu)$ were continuous in μ for fixed τ, since τ_t (for example) is an endogenously determined function of μ, τ will depend on μ as a measurable function which is generally not continuous. When the dependence of τ on μ is taken into account, $W_t(\alpha, \tau, \mu)$ depends measurably, but not continuously, on μ. As a result, the convergence of expressions such as $\int_{\mathcal{M}(A) \times A} W_t(\tilde{\alpha}, \tau, \tilde{\mu}) \mathbf{P}_{\mathcal{M}\xi}(d\tilde{\mu} \times d\tilde{\alpha}; y, \tau^k, \mu) \psi(d\mu)$ to $\int_{\mathcal{M}(A) \times A} W_t(\tilde{\alpha}, \tau, \tilde{\mu}) \mathbf{P}_{\mathcal{M}\xi}(d\tilde{\mu} \times \tilde{\alpha}; y, \tau, \mu) \psi(d\mu)$ depends on the assumption of norm continuity on measurable functions. Assumption 3 of uniform boundedness of the Radon-Nikodym derivative is made for similar reasons.

In this framework, the appropriate definition of equilibrium is:

Definition 4 *A strategy τ is a Markov equilibrium if τ is consistent with itself and for each t,*

$$\int_Y V_t(\alpha, a, \tau, \mu) \tau_t(dy; \mu) \psi(d\mu) \geq \int_Y V_t(\alpha, a, \tau, \mu) \tilde{\tau}_t(dy; \mu) \psi(d\mu),$$

for all $\tilde{\tau}$ consistent with τ.

Theorem 5 *The consistency and best response mappings satisfy the conditions of the Glicksberg Fan Theorem so that there exists a Markov equilibrium.*

5 Strategy specification

Earlier, when presenting the model we defined strategies as measurable functions from the aggregate shock history to distributions on characteristics and actions, and promised some further discussion of the formulation of strategies. We now describe an alternative approach to strategy formulation, that appears at first glance to have the appeal of simplicity.

One might think that rather than specify a period distributional strategy as a measurable function from the history of aggregate shocks to $\mathcal{M}(A \times A)$, that it would simplify the analysis to define the period t distributional strategy as a measure on $(\Theta^t \times A \times A)$. Such a formulation simplifies somewhat the choice of topology on strategies and appears at first glance to make the proof of existence of equilibrium easier (see Bergin and Bernhardt 1993). However, serious problems emerge when trying to interpret payoffs and transition probabilities defined on such strategies. For example, suppose that aggregate uncertainty concerns demand in a market which the agents are firms and that demand (θ) is uniformly distributed on the interval $[\theta_l, \theta_h]$, $\theta_h > \theta_l$. Consider period 1. The profit of firm α, supplying quantity a_θ when demand is θ, is given by $u_1(\alpha, a_\theta, \tau_1(\theta), \theta)$, where $\tau_1(\theta)$ represents the actions of other firms given that demand is θ. The expected profit to firm α in period 1 is $\int_\Theta u_1(\alpha, a_\theta, \tau_1(\theta), \theta) v(d\theta)$, with v the uniform distribution on $[\theta_l, \theta_h]$. If the distributional strategy is defined as a measure τ_1 on $(\Lambda \times A \times \Theta)$, calculation of expected profit of the firm still requires that we compute the conditional distribution over $\Lambda \times A$ for each $\theta \in \Theta$, $\tau_1(\theta_h)$, a measurable function from Θ to $\Lambda \times A$. Whichever way distributional strategies are defined, discussion of continuity of utility on distributional strategies involves interpreting the distributional strategy as a measurable function from Θ space to distributions on $(\Lambda \times A)$. This brings us back indirectly to the same formulation – of the strategy as a function from histories. Since the economic interpretation of the utility function requires the evaluation of distributional strategies on histories it is more useful to proceed directly with this formulation. By defining distributional strategies as measurable functions directly, i.e. $\tau_t \in \mathscr{F}(\Theta^t, \mathscr{M}(\Lambda \times A))$, we avoid the need to move back and forth between joint and conditional distributions. Similar comments apply concerning the transition probabilities.

A second virtue of our approach is that much weaker continuity assumptions are required to prove the existence of an equilibrium. We do not require continuity of the utility or transitions functions in the aggregate shock: neither $u_t(\alpha, a, m_t, \theta^t)$ nor $\mathbf{P}_{t+1}(m_t, \theta^t, y)$ need be continuous in θ^t. When the distributional strategy is defined as a joint distribution ($\tau_t \in \mathscr{M}(\Theta^t \times \Lambda \times A)$), continuity of both the utility and transition probabilities in θ^t is required to exploit the benefits of having τ_t defined as a distribution on $\Theta^t \times \Lambda \times A$.

6 Conclusion

As we observed in the introduction, there are many dynamic economic phenomena that can be fruitfully modelled as anonymous sequential games with aggregate uncertainty. In many economic environments, the stochastic evolution of both an individual agent's type and aggregate variables is key, as are their interactions. Our paper shows how these problems can be formulated, and provides conditions under which the economy's evolution has a Markov representation.

A classic environment in which such dynamics are key to our understanding is that of real business cycles. Davis, Haltiwanger and Schuh (1993), offer overwhelming evidence that different firms have very different experiences at the same moment in business cycles. They find, for instance, that in both upturns and downturns there is considerable entry and exit of plants and creation and destruction of jobs. As well, entry and exit of small firms varies less with the business cycle than the exit of large firms (which is highly countercyclical). Most job creation and job destruction in the economy is at a few firms – the big winners and losers; that is, job gains and losses are not evenly spread.

Macroeconomic models that do not incorporate both idiosyncratic and aggregate stochastic heterogeneity simply have no chance to explain these phenomena. Bergin and Bernhardt (1993) show how one can capture these phenomena within an anonymous sequential game which features a continuum support for the aggregate demand shocks, and provide Markov characterizations. Gouge and King (1994) use the formulation in a model that focuses on the labor side of these observations – human capital acquisition, job creation and job destruction.

Our framework can also be used to characterize asset market equilibria with incomplete markets, heterogeneous agents (e.g. with stochastically varying idiosyncratic endowments or preference shocks) and idiosyncratic and aggregate shocks to asset payoffs. DGMM provide conditions under which a Markov characterization of such a market equilibrium obtains in their stochastic games formulation, but one could provide related Markov characterizations in our anonymous sequential game formulation. To the extent that the technical conditions that we impose are less onerous, the anonymous sequential games formulation may allow better characterizations.

Another area in which our framework can be usefully employed is where firms which receive idiosyncratic cost shocks operate in an economy with aggregate inflation shocks. In such an environment one could study how and when firms adjust prices when price adjustment is costly. Important issues to consider include: how will inflation shocks be incorporated into prices? how does the mix between aggregate inflation shocks and idiosyncratic cost shocks affect price adjustment decisions? how will the pattern of price adjustment vary across markets? how will price adjustments be correlated within and across markets? how will the nature of competition within a market affect the responsiveness of a firm's pricing decisions to aggregate inflation shocks?, etc.

Thus, the anonymous sequential game with aggregate uncertainty formulation can be flexibly employed to model a host of economic phenomena. Our paper offers a very general equilibrium existence argument for these models, and provides conditions under which the economy has a Markov characterization. The researcher can then use the additional structure that he places on the economy to tease out further characterizations.

7 Appendix

We first give two lemmas. The first asserts that "continuity is preserved" through integration while the second lemma shows that continuity and norm continuity are "preserved under maximization". We use these lemmas to establish the existence of appropriate value functions. The value functions are then used to define best response mappings, leading to the equilibrium existence proof.

Lemma 1 *Let X, Y and Z be compact metric spaces. Let $\pi(x, y, z)$ be continuous in $(x, y, z) \in X \times Y \times Z$, and $\mathbf{P}(\bullet; y, z)$: $Y \times Z \to \mathcal{M}(X)$ be continuous in (y, z), so that $(y_k, z_k) \to (y, z)$ implies that the sequence of measures $\mathbf{P}(\bullet; y_k, z_k)$ converges weak* to $\mathbf{P}(\bullet; y, z)$. Finally, let $\mathbf{Q}(\bullet; z)$: $Z \to \mathcal{M}(Y)$ be weak* continuous in z. Then $z_k \to z$ implies* *that*

$$\int_Y \int_X \pi(x, y, z_k)\mathbf{P}(dx; y, z_k)\mathbf{Q}(dy; z_k) \to \int_Y \int_X \pi(x, y, z)\mathbf{P}(dz; y, z)\mathbf{Q}(dy; z).$$

Proof: Let $\gamma(y,z) = \int_X \pi(x,y,z)\mathbf{P}(dx; y, z)$ and note that $\gamma(y,z)$ is continuous. To see this, let $w = (y, z)$ and consider a sequence $w_k \to w$. Then

$$\begin{aligned}|\gamma(w_k) - \gamma(w)| &= \left| \int \pi(x, w_k)\mathbf{P}(dx; w_k) - \int \pi(x, w)\mathbf{P}(dx; w) \right| \\ &\leq \left| \int \pi(x, w_k)\mathbf{P}(dx; w_k) - \int \pi(x, w)\mathbf{P}(dx; w_k) \right| \\ &+ \left| \int \pi(x, w)\mathbf{P}(dx; w_k) - \int \pi(x, w)\mathbf{P}(dx; w) \right|.\end{aligned}$$

Since $|\int \pi(x, w_k)\mathbf{P}(dx; w_k) - \int \pi(x, w)\mathbf{P}(dx; w_k)| \leq \int |\pi(x, w_k) - \int \pi(x, w)| \mathbf{P}(dx; w_k)$ and π is uniformly continuous on $X \times W$ ($X \times W$ is a compact metric space), then given $\varepsilon > 0$, $\exists \bar{k}$ such that $k \geq \bar{k}$ implies that $|\pi(x, w_k) - \pi(x, w)| \leq \varepsilon$, for all x. Thus,

$$\left| \int \pi(x, w_k)\mathbf{P}(dx; w_k) - \int \pi(x, w)\mathbf{P}(dx; w_k) \right| \to 0.$$

Since π is continuous on (x, w) and $\mathbf{P}(\bullet; w_k)$ converges weakly to $P(\bullet; w)$,

$$\left| \int \pi(x, w)\mathbf{P}(dx; w_k) - \int \pi(x, w)\mathbf{P}(dx; w) \right| \to 0.$$

Thus, γ is continuous on $W = Y \times Z$. Since W is a compact metric space, γ is uniformly continuous on W.

Now, $\int_Y \int_X \pi(x, y, z_k)\mathbf{P}(dx; y, z_k)\mathbf{Q}(dy; z_k) = \int_Y \gamma(y, z_k)\mathbf{Q}(dy; z_k)$, so that using the uniform continuity of γ and *weak** convergence of $\mathbf{Q}(\bullet; z_k)$ to $\mathbf{Q}(\bullet; z)$, the same argument as above gives

$$\int_Y \gamma(y, z_k)\mathbf{Q}(dy; z_k) \to \int_Y \gamma(y, z)\mathbf{Q}(dy; z) = \int_Y \int_X \pi(x, y, z)\mathbf{P}(dz; y, z)\mathbf{Q}(dy; z),$$

which completes the proof. ■

For the next lemma, we need the following notation. Let $(\Omega, \mathscr{D}, \mu)$ be a given probability space and $\mathscr{M}(Y)$ the set of measures on $Y = A \times A$. Let $\mathscr{F}(\Omega, \mathscr{M}(Y))$ denote the set of measurable functions from Ω to $\mathscr{M}(Y)$. A sequence of measures $\{\tau^k\}$ in $\mathscr{F}(\Omega, \mathscr{M}(Y))$ converges to a measure τ if and only if

$$\int_\Omega \int_Y f(y)\tau^k(dy; \omega)g(\omega)\mu(d\omega) \to \int_\Omega \int_Y f(y)\tau(dy; \omega)g(\omega)\mu(d\omega),$$

$$\forall f \in \mathscr{C}(Y), g \in \mathbf{L}_1(\Omega, \mathscr{R}, \mu).$$

With this notation, we have:

Lemma 2 *Let* $r: Y \times \mathcal{M}(Y) \times \Omega \to \mathcal{R}$. *With a mild abuse of notation, write* $r(y, \tau, \omega)$ *to denote* $r(y, \tau(\omega), \omega)$ *(thus extending* r *to* $Y \times \mathcal{F}(\Omega, \mathcal{M}(Y)) \times \Omega$*). Let* $r(y, \tau, \omega)$ *be continuous on* Y *and norm continuous with respect to* τ: *as* $\tau^k \to \tau$, $\int_\Omega sup_y |r(y, \tau^k, \omega) - r(y, \tau, \omega)| \, \mu(d\omega) \to 0$. *Then* $s(\alpha, \tau, \omega) = max_a r(\alpha, a, \tau, \omega)$ *is continuous in* α *and norm continuous in* τ: $\tau^k \to \tau$ *implies* $\int_\Omega sup_\alpha |s(\alpha, \tau^k, \omega) - s(\alpha, \tau, \omega)| \mu(d\omega) \to 0$.

Proof: Continuity of s in α is clear. To consider norm continuity of s in τ, let $\tau^k \to \tau$. Since r is norm continuous, given $\varepsilon > 0, \exists \bar{k}$, such that $k \geq \bar{k}$ implies

$$\int_\Omega sup_y |r(\alpha, a, \tau^k, \omega) - r(\alpha, a, \tau, \omega)| \mu(d\omega) \leq \varepsilon.$$

Let

$$\Omega_k(\beta\varepsilon) = \{\omega \mid |r(\alpha, a, \tau^k, \omega) - r(\alpha, a, \tau, \omega)| \geq \beta\varepsilon\}.$$

Then

$$\begin{aligned}\varepsilon &\geq \int_\Omega sup_y |r(\alpha, a, \tau^k, \omega) - r(\alpha, a, \tau, \omega)| \mu(d\omega) \\ &\geq \int_{\Omega_k(\beta\varepsilon)} sup_y |r(\alpha, a, \tau^k, \omega) - r(\alpha, a, \tau, \omega)| \mu(d\omega) \geq \beta\varepsilon\mu(\Omega_k(\beta\varepsilon)).\end{aligned}$$

Thus, $1 \geq \beta\mu(\Omega_k(\beta\varepsilon))$, and setting $\beta = 1/\sqrt{\varepsilon}$ gives $\sqrt{\varepsilon} \geq \mu(\Omega_k(\sqrt{\varepsilon}))$. Let $a^k(\alpha, \omega)$ maximize $r(\alpha, a, \tau^k, \omega)$ and $a(\alpha, \omega)$ maximize $r(\alpha, a, \tau, \omega)$. On $\Omega_k(\sqrt{\varepsilon})^c$, $\forall \alpha$

$$r(\alpha, a^k(\alpha, \omega), \tau^k, \omega) \geq r(\alpha, a(\alpha, \omega), \tau^k, \omega) \geq r(\alpha, a(\alpha, \omega), \tau, \omega) - \sqrt{\varepsilon}.$$

The first inequality follows since $a^k(\alpha, \omega)$ is a maximizer of $r(\alpha, a, \tau^k, \omega)$ and the second inequality follows since $\omega \in \Omega_k(\sqrt{\varepsilon})^c$. Similarly,

$$r(\alpha, a^k(\alpha, \omega), \tau^k, \omega) \leq r(\alpha, a^k(\alpha, \omega), \tau, \omega) + \sqrt{\varepsilon} \leq r(\alpha, a(\alpha, \omega), \tau, \omega) + \sqrt{\varepsilon}.$$

The first inequality follows since $\omega \in \Omega_k(\sqrt{\varepsilon})^c$ and the second follows since $a(\alpha, \omega)$ is a maximizer of $r(\alpha, a, \tau, \omega)$. Consequently, $\forall \alpha$, $|s(\alpha, \tau^k, \omega) - s(\alpha, \tau, \omega)| \leq \sqrt{\varepsilon}, \omega \in \Omega_k(\sqrt{\varepsilon})^c$. Thus,

$$sup_\alpha |s(\alpha, \tau^k, \omega) - s(\alpha, \tau, \omega)| \leq \sqrt{\varepsilon}, \omega \in \Omega_k(\sqrt{\varepsilon})^c.$$

Therefore,

$$\begin{aligned}&\int_\Omega sup_\alpha |s(\alpha, \tau^k, \omega) - s(\alpha, \tau, \omega)| \mu(d\omega) \\ &\quad = \int_{\Omega_k(\sqrt{\varepsilon})} sup_\alpha |s(\alpha, \tau^k, \omega) - s(\alpha, \tau, \omega)| \mu(d\omega) \\ &\qquad + \int_{\Omega_k(\sqrt{\varepsilon})^c} sup_\alpha |s(\alpha, \tau^k, \omega) - s(\alpha, \tau, \omega)| \mu(d\omega).\end{aligned}$$

The latter expression is bounded above by $2K\mu(\Omega_k(\sqrt{\varepsilon})) + \sqrt{\varepsilon}\mu(\Omega_k(\sqrt{\varepsilon})^c) \leq 2K\sqrt{\varepsilon} + \sqrt{\varepsilon} = [2K+1]\sqrt{\varepsilon}$. ■

Theorem 1 *For each t, there exist value functions $V_t(\alpha, a, \tau, \theta^t)$ and $W_t(\alpha, \tau, \theta^t)$, which are continuous in (α, a) and α respectively, norm continuous in τ and satisfy $W_t(\alpha, \tau, \theta^t) = max_a V_t(\alpha, a, \tau, \theta^t)$.*

Proof: Consider an n-period truncation of the game. Trivially in period n, $V_n^n(\alpha, a, \tau, \theta^n) = u_n(\alpha, a, \tau_n, \theta_n)$, which, by assumption, is continuous in (α, a) and norm continuous in τ, as is $W_n^n(\alpha, \tau, \theta^n) = max_a u_n(\alpha, a, \tau_n, \theta_n)$, by lemma 2. Now define

$$V_{n-1}^n(\alpha, a, \tau, \theta^{n-1}) = u_{n-1}(\alpha, a, \tau_{n-1}, \theta_{n-1}) + \delta_n \int_\Theta \int_A W_n^n(\xi, \tau, \theta^n)\mathbf{P}_n(d\xi, \tau_{n-1}, \theta^{n-1}, \alpha, a)v_n(d\theta^n|\theta^{n-1}).$$

$V_{n-1}^n(\alpha, a, \tau, \theta^{n-1})$ satisfies: (1) $V_{n-1}^n(\alpha, a, \tau, \theta^n)$ is continuous in (α, a) and (2) if $\tau^k \to \tau$, then $\int_\Theta sup_{\alpha,a}|V_{n-1}^n(\alpha, a, \tau^k, \theta^{n-1}) - V_{n-1}^n(\alpha, a, \tau, \theta^{n-1})|v_n(d\theta^{n-1}) \to 0$. Continuity in (α, a) follows from lemma 1, treating $(\tau_{n-1}, \theta^{n-1})$ as parameters of u_{n-1} and $\mathbf{P}_n$, respectively. Norm continuity in τ can be seen by separating current and future components of expected payoffs, so that, using abbreviated notation (writing $\mathbf{P}_n^k$ for $\mathbf{P}_{\xi_n}^k(d\xi, \tau_{n-1}^k, \theta^{n-1}, \alpha, a)$, etc.),

$$\int_\Theta sup_{(\alpha,a)}|V_{n-1}^n(\alpha, a, \tau^k, \theta^{n-1}) - V_{n-1}^n(\alpha, a, \tau, \theta^{n-1})|v_n(d\theta^{n-1})$$

becomes

$$\int_{\Theta^{n-1}} sup_{(\alpha,a)}\left|u_{n-1}^k - u_{n-1} + \delta_n \int_\Theta \int_A W_n^{nk}\mathbf{P}_n^k v_n(d\theta^n|\theta^{n-1}) - \delta_n \int_\Theta \int_A W_n^n \mathbf{P}_n v_n(d\theta^n|\theta^{n-1})\right| v_{n-1}(d\theta^{n-1})$$

$$\leq \int_{\Theta^{n-1}} sup_{(\alpha,a)}|u_{n-1}^k - u_{n-1}|v(d\theta^{n-1}) + \delta_n \int_{\Theta^{n-1}} sup_{(\alpha,a)}\left|\int_\Theta \int_A W_n^{nk}\mathbf{P}_n^k v_n(d\theta^n|\theta^{n-1}) - \int_\Theta \int_A W_n^n \mathbf{P}_n v_n(d\theta^n|\theta^{n-1})\right| v_{n-1}(d\theta^{n-1}).$$

The first term on the right hand side converges to 0 as $\tau^k \to \tau$, because u_{n-1} is norm continuous. For the second term consider

$$\int_{\Theta^{n-1}} sup_{(\alpha,a)}\left|\int_\Theta \int_A W_n^{nk}\mathbf{P}_n^k v_n(d\theta^n|\theta^{n-1}) - \int_\Theta \int_A W_n^n \mathbf{P}_n v_n(d\theta^n|\theta^{n-1})\right| v_{n-1}(d\theta^{n-1})$$

$$\leq \int_{\Theta^{n-1}} \int_\Theta sup_{(\alpha,a)}\left|\int_A W_n^{nk}\mathbf{P}_n^k - \int_A W_n^n \mathbf{P}_n\right| v_n(d\theta^n|\theta^{n-1})v_{n-1}(d\theta^{n-1})$$

$$= \int_{\Theta^n} sup_{(\alpha,a)} \left| \int_A W_n^{nk} \mathbf{P}_n^k - \int_A W_n^n \mathbf{P}_n \right| v_n(d\theta^n).$$

The last expression is no greater than:

$$\int_{\Theta^n} sup_{(\alpha,a)} \left| \int_A W_n^{nk} \mathbf{P}_n^k - \int_A W_n^n \mathbf{P}_n^k \right| v_n(d\theta^n)$$
$$+ \int_{\Theta^n} sup_{(\alpha,a)} \left| \int_A W_n^n \mathbf{P}_n^k - \int_A W_n^n \mathbf{P}_n \right| v_n(d\theta^n).$$

The first term is bounded above by $\int_{\Theta^n} sup_\alpha | \int_A W_n^{nk} - W_n^n | v_n(d\theta^n)$, which converges to 0, from norm continuity of W_n^n. The second term converges to 0, by norm continuity of $\mathbf{P}_n^k(d\xi, \tau_{n-1}, \theta^{n-1}, \alpha, a)$ in τ. $W_{n-1}^n(\alpha, \tau, \theta^{n-1})$ is defined from $V_{n-1}^n(\alpha, a, \tau, \theta^{n-1})$ and, as before, is continuous in α. Norm continuity of $W_{n-1}^n(\alpha, \tau, \theta^{n-1})$ in τ follows from lemma 2.

Proceed inductively in this way to define $V_t^n(\alpha, a, \tau, \theta^t)$ and $W_t^n(\alpha, \tau, \theta^t)$ for $1 \le t \le n$. The discussion above defines the recursion for fixed n and shows that for any t, $1 \le t \le n$, that both $V_t^n(\alpha, a, \tau, \theta^t)(\equiv V_{n-(n-t)}^n(\alpha, a, \tau, \theta^{n-(n-t)}))$ and $W_t^n(\alpha, \tau, \theta^t)$ are continuous functions of (α, a) and α respectively, and that both are norm continuous in τ. To conclude we show that the following limits exist for each j and are continuous functions of (α, a) and α respectively:

$$lim_{n\to\infty} V_j^n(\alpha, a, \tau, \theta^j) = V_j(\alpha, a, \tau, \theta^j) \text{ and } lim_{n\to\infty} W_j^n(\alpha, \tau, \theta^j) = W_j(\alpha, \tau, \theta^j).$$

Taking $n > j$, observe that each of the functions $V_j^n(\alpha, a, \tau, \theta^j)$ and $W_j^n(\alpha, \tau, \theta^j)$ is increasing in n, and that

$$0 \le V_j^{n+s}(\alpha, a, \tau, \theta^j) - V_j^n(\alpha, a, \tau, \theta^j) \le \sum_{r=1}^{s} (\times_{h=j}^{n+r-1} \delta_{j+h})K \le [\delta^{n-j+1}/(1-\delta)]K,$$

$$0 \le W_j^{n+s}(\alpha, \tau, \theta^j) - W_j^n(\alpha, \tau, \theta^j) \le \sum_{r=1}^{s} (\times_{h=j}^{n+r-1} \delta_{j+h})K \le [\delta^{n-j+1}/(1-\delta)]K.$$

Therefore $V_j^n(\alpha, a, \tau, \theta^j)$ and $W_j^n(\alpha, \tau, \theta^j)$ are Cauchy sequences in n. Since $V_j^n(\alpha, a, \tau, \theta^j)$ and $W_j^n(\alpha, \tau, \theta^j)$ are continuous in (α, a) and α respectively and are both norm continuous in τ, the limits $V_t(\alpha, a, \tau, \theta^t) = lim_n V_t^n(\alpha, a, \tau, \theta^t)$ and $W_t(\alpha, \tau, \theta^t) = lim_n W_t^n(\alpha, \tau, \theta^t)$ inherit these properties also. ■

Theorem 2 *The correspondence* $\mathbf{C}(\tau) \equiv \times_{t=1}^{\infty} \mathbf{C}_t(\tau)$ *is non-empty, upper-hemicontinuous and convex-valued.*

Proof: Let $\tilde{\tau}^n = (\tilde{\tau}_1^n, \ldots, \tilde{\tau}_t^n, \ldots)$ and $\tau^n = (\tau_1^n, \ldots, \tau_t^n, \ldots)$ where $\tilde{\tau}_t^n, \tau_t^n \in \mathscr{F}(\Theta^t, \mathscr{M}(A \times A))$ with $\tilde{\tau}^n \in C(\tau^n)$, $\tilde{\tau}_t^n \to \tilde{\tau}_t$, $\tau_t^n \to \tau_t$, and

$$\int_{\Theta^t} \tilde{\tau}_t^n(f, A; \theta^t) g(\theta^t) v_t(d\theta^t) = \int_{\Theta^t} \int_Y \mathbf{P}_t(f, \tau_{t-1}^n, \theta^{t-1}, y) \tau_{t-1}^n(dy; \theta^{t-1}) g(\theta^t) v_t(d\theta^t).$$

Then,

$$\int_{\Theta^t} \tilde{\tau}_t(f, A; \theta^t) g(\theta^t) v_t(d\theta^t) = \int_{\Theta^t} \int_Y \mathbf{P}_t(f, \tau_{t-1}, \theta^{t-1}, y) \tau_{t-1}(dy; \theta^{t-1}) g(\theta^t) v_t(d\theta^t).$$

To see this, note directly from the topology on τ that

$$\int_{\Theta^t} \tilde{\tau}_t^n(f, A; \theta^t) g(\theta^t) v_t(d\theta^t) \to \int_{\Theta^t} \tilde{\tau}_t(f, A; \theta^t) g(\theta^t) v_t(d\theta^t).$$

Now consider the right hand side. Abbreviate $\mathbf{P}_t(f, \tau_{t-1}^n, \theta^{t-1}, y)$ by $\mathbf{P}_t^n(\theta^{t-1}, y)$ and $\mathbf{P}_t(f, \tau_{t-1}, \theta^{t-1}, y)$ by $\mathbf{P}_t(\theta^{t-1}, y)$. Then

$$\begin{aligned}
&\left| \int_{\Theta^t} \int_Y \mathbf{P}_t^n(\theta^{t-1}, y) \tau_{t-1}^n(dy; \theta^{t-1}) g dv_t - \int_{\Theta^t} \int_Y \mathbf{P}_t(\theta^{t-1}, y) \tau_{t-1}(dy; \theta^{t-1}) g dv_t \right| \\
&\quad \leq \left| \int_{\Theta^t} \int_Y \mathbf{P}_t^n(\theta^{t-1}, y) \tau_{t-1}^n(dy; \theta^{t-1}) g dv_t - \int_{\Theta^t} \int_Y \mathbf{P}_t(\theta^{t-1}, y) \tau_{t-1}^n(dy; \theta^{t-1}) g dv_t \right| \\
&\quad + \left| \int_{\Theta^t} \int_Y \mathbf{P}_t(\theta^{t-1}, y) \tau_{t-1}^n(dy; \theta^{t-1}) g dv_t - \int_{\Theta^t} \int_Y \mathbf{P}_t(\theta^{t-1}, y) \tau_{t-1}(dy; \theta^{t-1}) g dv_t \right|.
\end{aligned}$$

The first term on the right of the inequality is less than or equal to

$$\begin{aligned}
&\int_{\Theta^t} \int_Y \| \mathbf{P}_t^n(\theta^{t-1}, y) - \mathbf{P}_t(\theta^{t-1}, y) \|_y \tau_{t-1}^n(dy; \theta^{t-1}) g dv_t \\
&\quad = \int_{\Theta^t} \| \mathbf{P}_t^n(\theta^{t-1}, y) - \mathbf{P}_t(\theta^{t-1}, y) \|_y g dv_t,
\end{aligned}$$

and the norm continuity condition on $\mathbf{P}_t$ implies that this goes to zero. The second term converges to zero from the topology on τ. Convexity follows since the restrictions are linear. It remains to show that non-emptyness is also satisfied. To see this, given $\mu_1(\bullet)$, the initial measure on A, and given the measure v_1 on Θ, let h be a measurable function from $A \times \Theta$ to A. Define a measure φ on $A \times A \times \Theta$ according to the property that $\varphi(X \times Z) = \mu_1 \otimes v_1(h^{-1}(X) \cap Z)$ for any measurable sets X and Z in A and $A \times \Theta$ respectively. (Interpret φ as the unique extension from such measurable rectangles). Let $\hat{\tau}_1(\bullet, \bullet; \theta_1) = \varphi(\bullet, \bullet; \theta_1)$, where $\varphi(\bullet, \bullet; \theta_1)$ is the conditional distribution of φ on $A \times A$, given θ_1. Note that $\varphi(A \times Z) = \mu_1 \otimes v_1(Z)$ so that $\hat{\tau}_1(\bullet, A; \theta_1) = \mu_1 \otimes v_1(\bullet; \theta_1) = \mu_1(\bullet)$. For $t \geq 2$, a similar discussion applies. View $\int_Y \mathbf{P}_t(f, \tau_{t-1}, \theta^{t-1}, y) \tau_{t-1}(dy; \theta^{t-1})$ as a conditional distribution on A given θ^{t-1}. Let $\mathbf{Q}$ be the joint distribution on $A \times \Theta^t$ determined by $\int_Y \mathbf{P}_t(f, \tau_{t-1}, \theta^{t-1}, y) \tau_{t-1}(dy; \theta^{t-1})$ and v_t. As before, let h be a measurable function from $A \times \Theta^t$ to A. Define a measure on $A \times A \times \Theta^t$, φ, determined on rectangles $X \times Z$, where X and Z are measurable subsets of A and $A \times \Theta^t$ respectively, by $\varphi(X \times Z) = \mathbf{Q}(h^{-1}(X) \cap Z)$. Let $\tilde{\tau}_t(\bullet, \bullet; \theta^t) = \varphi(\bullet, \bullet; \theta^t)$. Lastly, note that $\tilde{\tau}_t(\bullet, A; \theta^t) = \varphi(\bullet, A; \theta^t) = \mathbf{Q}(h^{-1}(A) \cap \bullet; \theta^t) = \mathbf{Q}(\bullet; \theta^t) = \int_Y \mathbf{P}_{\xi t}(\bullet, \tau_{t-1}, \theta^{t-1}, y) \tau_{t-1}(dy; \theta^{t-1})$. ■

Theorem 3 *The correspondence* **B** *satisfies the conditions of the Glicksberg Fan theorem and hence has a fixed point, which is an equilibrium of the game.*

Proof: To see that **B** is convex-valued, recall that **C** is convex-valued and the additional constraints on $\hat{\tau}$ in the definition of **B** are defined by linear inequalities, so that **B** is convex-valued. Next **B** is non-empty since for any t, $\mathbf{C}_t(\tau)$ is closed and non-empty (in view of theorem 2) and $\int_{\Theta^t}\int_Y V_t(\alpha, a, \tau, \theta^t)\tau_t(dy;\theta^t)\nu_t(d\theta^t)$ is continuous in τ. It remains to show that **B** is upper-hemicontinuous. We prove this in two steps.

Step 1: We first show that $\hat{\tau}\in\mathbf{B}(\tau)$ if and only if,

$$\int_{\Theta^t}\int_Y V_t(\alpha, a, \tau, \theta^t)\hat{\tau}_t(dy;\theta^t)\nu_t(d\theta^t) = \int_{\Theta^t}\int_A W_t(\alpha, \tau, \theta^t)\lambda_t(d\alpha;\theta^t)\nu_t(d\theta^t),$$

where $\lambda_t(\bullet;\theta^t)$ is the (unique) marginal distribution on characteristics, A, determined by any distribution consistent with τ.[10] To see this, consider the correspondence

$$\psi(\alpha,\theta^t) \equiv \{a \mid max_a V_t(\alpha, a, \tau, \theta^t) \le V_t(\alpha, a, \tau, \theta^t)\} \equiv \{a \mid W_t(\alpha, \tau, \theta^t) \le V_t(\alpha, a, \tau, \theta^t)\}.$$

Denote the graph of ψ by $\mathscr{G}_\psi$ and observe that $\mathscr{G}_\psi \in \mathscr{B}_A \times \mathscr{B}_A \times \mathscr{B}^t_\Theta$ (ψ has a measurable graph) since

$$\mathscr{G}_\psi = \{(\alpha, a, \theta^t) \mid a\in\psi(\alpha,\theta^t)\} = \{(\alpha, a, \theta^t) \mid W_t(\alpha,\tau,\theta^t) \le V_t(\alpha, a, \tau, \theta^t)\}.$$

Denote by $\lambda\otimes\nu_t$ the measure on $A\times\Theta^t$ determined by λ and ν_t. Viewing ψ as a correspondence from $(A\times\Theta^t, \mathscr{B}_A\times\mathscr{B}^t, \lambda\otimes\nu_t)$, there is a measurable selection $h: A\times\Theta^t \to A$, with $h(\alpha,\theta^t)\in\psi(\alpha,\theta^t)$ almost everywhere $\lambda\otimes\nu_t$ since $\mathscr{G}_\psi\in\mathscr{B}_A\times\mathscr{B}_A\times\mathscr{B}^t$ (using the measurable selection theorem). Thus $V_t(\alpha, h(\alpha,\theta^t), \tau, \theta^t) = W_t(\alpha,\tau,\theta^t)$, almost everywhere $\lambda\otimes\nu_t$. Now, define a distribution on $A\times A\times\Theta^t$ by $\varphi(W\times Z) = \lambda\otimes\nu_t(h^{-1}(W)\cap Z), \forall\, W\in\mathscr{B}_A$ and $Z\in\mathscr{B}_A\times\mathscr{B}^t$. Observe that $\varphi(A\times Z) = (\lambda\otimes\nu_t)(Z)$ so that ν_t almost everywhere θ^t, $\varphi(A,\bullet;\theta^t) = \lambda(\bullet;\theta^t)$. Thus, φ has two key features: for almost all θ^t, the conditional distribution on $A\times A$, given θ^t has support on the best response mapping (h) from A to A and the marginal on A has the required consistency property. Define $\tau^*_t : \tau^*_t(\bullet,\bullet;\theta^t) = \varphi(\bullet,\bullet;\theta^t)$. Then $\tau^*_t\in\mathbf{C}_t(\tau)$ and,

$$\int_{\Theta^t}\int_Y V_t(\alpha, a, \tau, \theta^t)\tau^*_t(dy;\theta^t)\nu_t(d\theta^t)$$
$$= \int_{\Theta^t}\int_A V_t(\alpha, h(\alpha,\theta^t), \theta^t)\lambda_t(d\alpha;\theta^t)\nu_t(d\theta^t) \ge \int_{\Theta^t}\int_A W_t(\alpha,\tau,\theta^t)\lambda_t(d\alpha;\theta^t)\nu_t(d\theta^t).$$

However, since $W_t(\alpha,\tau,\theta^t) \ge V_t(\alpha, a, \tau, \theta^t), \forall(\alpha, a)\in A\times A$, we have

$$\int_{\Theta^t}\int_Y V_t(\alpha, a, \tau, \theta^t)\tau^*_t(dy;\theta^t)\nu_t(d\theta^t) = \int_{\Theta^t}\int_A W_t(\alpha,\tau,\theta^t)\lambda_t(d\alpha;\theta^t)\nu_t(d\theta^t).$$

Step 2: Now, let $\tau^n\to\tau$ and suppose that $\tilde{\tau}^n\in\mathbf{B}(\tau^n)$, with $\tilde{\tau}^n\to\tilde{\tau}$. It is necessary to show that $\tilde{\tau}\in\mathbf{B}(\tau)$. Recall $W_t(\alpha,\tau,\theta^t)$ is norm continuous:

$$\int_{\Theta^t} sup_\alpha |W_t(\alpha,\tau^n,\theta^t) - W_t(\alpha,\tau,\theta^t)|\,\nu(d\theta^t)\to 0.$$

Let $\lambda^n_t(\bullet;\theta^t)$ be the distribution on A determined by $\mathbf{C}_t(\tau^n)$, so that if $\hat{\tau}^n_t\in\mathbf{C}(\tau^n)$, then $\hat{\tau}^n_t(\bullet, A;\theta^t) = \lambda^n_t(\bullet;\theta^t)$, ν_t almost everywhere. Norm continuity of $W_t(\alpha,\tau,\theta^t)$ implies

[10] Recall that each distribution in $\mathbf{C}_t(\tau)$ must have the same marginal distribution on A: if $\hat{\tau}_t$, $\tilde{\tau}_t\in\mathbf{C}_t(\tau)$, then $\tau_t(\bullet, A;\theta^t) = \tau_t(\bullet, A;\theta^t)$, ν_t almost everywhere θ^t. We denote this distribution on A by $\lambda_t(\bullet;\theta^t)$.

that

$$\int_{\Theta^t}\int_A |W_t(\alpha,\tau^n,\theta^t)-W_t(\alpha,\tau,\theta^t)|\lambda_t^n(d\alpha;\theta^t)v(d\theta^t)\to 0.$$

Hence,

$$\left|\int_{\Theta^t}\int_A W_t(\alpha,\tau^n,\theta^t)\lambda_t^n(d\alpha;\theta^t)v(d\theta^t)-\int_{\Theta^t}\int_A W_t(\alpha,\tau,\theta^t)\lambda_t^n(d\alpha;\theta^t)v(d\theta^t)\right|\to 0.$$

Now in the topology given on measures, let λ_t be the limit of λ_t^n and note that $\hat{\tau}_t\in\mathbf{C}(\tau)$, so $\hat{\tau}_t(A,\bullet;\theta^t)=\lambda_t(\bullet;\theta^t)$, v_t almost everywhere. Observe also that $W_t(\alpha,\tau,\theta^t)$ is continuous in α (since $W_t(\alpha,\tau,\theta^t)=max_a V_t(\alpha,a,\tau,\theta^t)$ and $V_t(\alpha,a,\tau,\theta^t)$ is continuous in (α,a), so that $W_t(\alpha,\tau,\theta^t)$ is continuous in α for each θ^t). Thus

$$\int_{\Theta^t}\int_A W_t(\alpha,\tau,\theta^t)\lambda_t^n(d\alpha;\theta^t)v(d\theta^t)\to\int_{\Theta^t}\int_A W_t(\alpha,\tau,\theta^t)\lambda_t(d\alpha;\theta^t)v(d\theta^t),$$

so that

$$\int_{\Theta^t}\int_A W_t(\alpha,\tau^n,\theta^t)\lambda_t^n(d\alpha;\theta^t)v(d\theta^t)\to\int_{\Theta^t}\int_A W_t(\alpha,\tau,\theta^t)\lambda_t(d\alpha;\theta^t)v(d\theta^t).$$

Now recall that since $\tilde{\tau}^n\in\mathbf{B}(\tau^n)$, by step 1:

$$\int_{\Theta^t}\int_A V_t(\alpha,a,\tau^n,\theta^t)\tilde{\tau}_t^n(d\alpha\times da;\theta^t)v(d\theta^t)=\int_{\Theta^t}\int_A W_t(\alpha,\tau^n,\theta^t)\lambda_t^n(d\alpha;\theta^t)v(d\theta^t).$$

Since $\tilde{\tau}^n\to\tilde{\tau}$ as $\tau^n\to\tau$, using the norm continuity of $V_t(\alpha,a,\tau^n,\theta^t)$ we find that

$$\int_{\Theta^t}\int_A V_t(\alpha,a,\tau^n,\theta^t)\tilde{\tau}_t^n(d\alpha\times da;\theta^t)v(d\theta^t)\to\int_{\Theta^t}\int_A V_t(\alpha,a,\tau,\theta^t)\tilde{\tau}_t(d\alpha\times da;\theta^t)v(d\theta^t).$$

Therefore,

$$\int_{\Theta^t}\int_A V_t(\alpha,a,\tau,\theta^t)\tilde{\tau}_t(d\alpha\times da;\theta^t)v(d\theta^t)=\int_{\Theta^t}\int_A W_t(\alpha,\tau,\theta^t)\lambda_t(d\alpha;\theta^t)v(d\theta^t).$$

Thus, $\tilde{\tau}\in\mathbf{B}(\tau)$, so that $\mathbf{B}$ is upper-hemicontinuous. Therefore $\mathbf{B}$ is convex-valued, non-empty and upper-hemicontinuous and so has a fixed point. ■

Theorem 4 *Given an equilibrium τ of the game with initial characteristics distribution μ and initial state θ, there is a Markov equilibrium, $\bar{\tau}$, such that the first period payoff to each agent is unchanged: the expected payoff to α is the same under $\bar{\tau}$ as τ.*

Proof: In view of the following facts:

1. $\mathbf{E}$ is upper-hemicontinuous,
2. W_1 is norm continuous in τ, and continuous in (α,θ),
3. $\mathcal{M}_\infty$ is metrizable and compact,

it follows that φ is an upper-hemicontinuous correspondence into a complete separable space. Hence there is a *pointwise* measurable selection, $\tau^*,\tau^*(\mu,v,\theta)\in$

$\varphi(\mu, v, \theta)$, for all $(\mu, v, \theta) \in \mathbf{S}$. We use τ^* to construct the Markov equilibrium $\bar{\tau}$. Consider the first component of $\tau^*(\mu, v, \theta)$, $\tau_1^*(\mu, v, \theta)$. This is a measure on $A \times A$ which is optimal in the sense that at (μ, v, θ):

$$\tau_1^*(\mu, v, \theta)\{(\alpha, a) \mid V_1(\alpha, a, \tau^*, \theta) \geq W_1(\alpha, \tau^*, \theta)\} = 1.$$

To implement the strategy in period one, knowledge of (μ, v, θ) is required. Now given θ, let

$$\mu_2(\bullet \mid \theta) = \int_Y \mathbf{P}(\bullet, \tau_1^*, \theta, y)\tau_1^*(dy).$$

The measure $\mu_2(\bullet \mid \theta)$ is the second period distribution on characteristics. Given a realization of the aggregate shock in the second period, say θ', the expected payoff to agent α over the remainder of the game is: $W_2(\alpha, \tau^*, (\theta, \theta'))(= v|_{(\theta,\theta')}(\alpha))$. Considering τ^* and (θ, θ') fixed, $W(\alpha, \tau^*, (\theta, \theta'))$ is an element of $C(\Lambda))$, which we can write as $v_2(\alpha)$. Now observe that τ^* induces an equilibrium from period 2 on, for all "histories" except possibly a set of v measure 0. Thus, except for a set θ''s of v measure 0,

$$(\mu_2(\bullet \mid \theta), W_2(\alpha, \tau^*, (\theta, \theta')), \theta') \in \mathbf{S}.$$

Denote this "state" by (μ_2, v_2, θ'). Viewed as a subgame, the expected payoff to agent α is $v_2(\alpha)$. Note that this payoff is generated at this subgame by τ^*: $\tau_2^*(\bullet, \bullet \mid (\theta, \theta'))$. However, note that exactly the same payoff is obtained on this subgame if $(\tau_2^*, \tau_3^*, \ldots)$ is replaced by $\tau^*(\mu_2, v_2, \theta')$. For this reason, τ_1^* remains optimal and the stregegy obtained in this way is an equilibrium. Denote this strategy by $\tau^*(2) \in \mathscr{M}_\infty$ (given the initial θ) as

$$\tau^*(2) = (\tau_1^*(\mu, v, \theta), \tau^*[\mu_2(\bullet \mid \theta), W_2(\alpha, \tau^*, (\theta, \theta')), \theta']_{\theta \in \Theta}).$$

Thus, $\tau^*(2)$ is composed of $\tau_1^*(\mu, v, \theta)$ in the first period, and then τ^* is "restarted" in period 2: at the subgame reached by history θ', the "state" is $s_2 = (\mu_2(\bullet \mid \theta), W_2(\alpha, \tau^*, (\theta, \theta')), \theta')$ and this state is sustained by the strategy $\tau^*(s_2)$, at that subgame.

The important point about this construction is that τ_1^* is being applied at *period two* and this is the way in which Markov stationarity is introduced. Note that $\tau^*(2)$ induces an equilibrium on almost all histories, (θ, θ'), and gives the same continuation payoffs from the second period at each history as did τ^*. This ensures that τ_1^* is optimal at almost all θ in period one. Thus, the strategy $\tau^*(2)$ is also an equilibrium which gives the same first period payoff (v) as τ^*. First period "strategies", $\tau_1^*(\mu, v, \theta)$ are unchanged while second period strategies under $\tau^*(2)$ generate the same expected payoff there as did τ^*. The result of this construction is that the Markov property holds for the first and second period.

Now, replace the equilibrium strategy τ^* by the equilibrium strategy $\tau^*(2)$. This alters the evolution of the characteristics distribution and the valuation functions. In particular,

$$\mu_3(\bullet \mid \theta, \theta') = \int_Y \mathbf{P}(\bullet, \tau_2^*(2)[\mu_2(\bullet \mid \theta), W_2(\alpha, \tau_2^*(2), (\theta, \theta')), \theta'], \theta', y)$$
$$\times\, \tau_2^*(2)[\mu_2(\bullet \mid \theta), W_2(\alpha, \tau^*, (\theta, \theta'), \theta'](dy).$$

Similarly, there is a valuation function for period 3, $W_3(\alpha, \tau^*(2), (\theta, \theta', \tilde{\theta}))$. Here again, for a fixed history, $(\theta, \theta', \tilde{\theta})$, $W_3(\alpha, \tau^*(2), (\theta, \theta', \tilde{\theta})) \in C(A)$. Now, define $\tau^*(3)$

$$\tau^*(3) = (\tau_1^*(2), \tau_2^*(2), \tau^*[(\mu_3(\bullet|\theta, \theta'), W_2(\alpha, \tau^*, (\theta, \theta', \tilde{\theta})), \tilde{\theta}]_{(\theta,\theta') \in \Theta^2}).$$

As with $\tau^*(2)$, $\tau^*(3)$ is an equilibrium. Proceed in this way to define iteratively a sequence of equilibria $\tau^*(n)$ from $\tau^*(n-1)$ and observe that the sequence $\{\tau^*(n)\}_n$ converges, say to $\bar{\tau}$. Under $\bar{\tau}$ and the Markov distribution on Θ, the state variable $s = (\mu, v, \theta)$ evolves stochastically as a Markov chain. Schematically,

$$s_1 = (\mu, v|_\theta, \theta) \xrightarrow{\theta'} (\mu|_\theta, v|_{(\theta,\theta')}, \theta') = s_2 \xrightarrow{\tilde{\theta}} (\mu|_{\theta,\theta'}, v|_{(\theta,\theta',\tilde{\theta})}, \tilde{\theta})$$

$$= s_3 \xrightarrow{\theta^0} (\mu|_{(\theta,\theta',\tilde{\theta})}, v|_{(\theta,\theta',\tilde{\theta},\theta^0)}, \theta^0).$$

or alternatively,

$$s_1 = (\mu, v|_\theta, \theta) \xrightarrow{\theta'} (\mu|_\theta, v|_{(\theta,\theta')}, \theta') = s_2 = (\mu', v', \theta')$$

$$\xrightarrow{\tilde{\theta}} (\mu'|_{\theta'}, v'|_{\tilde{\theta}}, \tilde{\theta}) = (\hat{\mu}, \hat{v}, \tilde{\theta}) = s_3 \xrightarrow{\theta^0} (\hat{\mu}|_{\theta^0}, \hat{v}|_{\theta^0}, \theta^0).$$

The evolution of the states may be described as follows. With $\bar{\tau}$, given s_1, the distributional strategy at time 1 is $\tau_1^*(s_1)$. (Note that the first components of τ^* and $\bar{\tau}$ are related: $\tau_1^*(\mu, v, \theta) = \bar{\tau}_1(\mu, v, \theta)$, $\forall(\mu, v, \theta)$. At $t = 2$, the distributional strategy is $\tau_1^*(s_2)$, and at time $t, \tau_1^*(s_t)$. The influence of the θ sequence on strategies is only through the s variables, since given τ, s_t depends on $\theta^t = (\theta_1, \theta_2, \ldots, \theta_t)$. The behavior of $\bar{\tau}$ throughout the remainder of the game (from period t on) depends only on θ^t through s_t, so we can write the value function $W_t(\alpha, \tau, \theta^t)$ as $\bar{W}_t(\alpha, \tau, \theta_t, s_t(\theta^{t-1}))$. Note also that, since the environment is stationary, if $s_l(\theta^{l-1}) = s_t(\theta^{t-1})$, then $(\tau_1^*, s_l(\theta^{l-1}))$ and $(\tau_1^*, s_t(\theta^{t-1}))$ induce the same distribution over the state space in subsequent periods so that $\bar{W}_t(\alpha, \bar{\tau}, \theta, s_t(\theta^{t-1})) = \bar{W}_l(\alpha, \bar{\tau}, \theta, s_l(\theta^{l-1}))$. Consequently, we may write $\bar{W}(\alpha, \bar{\tau}, \theta, s_t(\theta^{t-1}))$ to denote the time t value function (without the time subscript).

A play of the game in this formulation may be described as follows. Fix an initial state $s = (\mu, v, \theta)$. At time $t = 1$ the distributional strategy $\tau_1^*(s)$ is played. Depending on the realization of the second period aggregate uncertainty variable, θ', a new state $s_2 = (\mu|_\theta, v|_{(\theta,\theta')}, \theta')$ is reached. The first component of $s_2 = (\mu|_\theta, v|_{(\theta,\theta')}, \theta')$ is equal to $\mu_2(\bullet|\theta) = \int_Y \mathbf{P}(\bullet, \tau_1^*, \theta, y)\tau_1^*(dy)$ and the second component is equal to $\bar{W}_2(\alpha, \bar{\tau}, \theta, s_2) = \bar{W}(\alpha, \theta, \bar{\tau}, s_2)$. For fixed s_2 and given $\bar{\tau}$, $\bar{W}(\alpha, a, \theta, \bar{\tau}, s_2) \in \mathbf{L}_1(\Theta, \mathscr{C}(A), v_1)$ and for fixed θ, $\mu_2(\bullet|\theta) \in \mathscr{M}(A)$. This completes the description of the Markov equilibrium. ■

Theorem 5 *There exists a Markov equilibrium.*

Proof: The proof follows essentially the same plan as the proof of theorem 3. This requires showing first that the consistency mapping is an upper-hemicontinuous correspondence and that there exist value functions for this case, analogous to those given in theorem 1. To define the value functions, follow theorem 1 and consider a game truncated to n periods. Given $\tau = \{\tau_t\}_{t=1}^\infty$, define $V_n^n(\alpha, a, \tau, \mu) = u(\alpha, a, \tau_n, \mu)$. Continuity in (α, a) and norm continuity in τ of $V_n^n(\alpha, a, \tau, \mu)$ follow directly since

$u(\alpha, a, \tau_n, \mu)$ has these properties. Let $W_n^n(\alpha, \tau, \mu) = max_a V_n^n(\alpha, a, \tau, \mu)$. $W_n^n(\alpha, \tau, \mu)$ is continuous in α and norm continuous in τ by lemma 1. Next, define

$$V_{n-1}^n(\alpha, a, \tau, \mu) = u(\alpha, a, \tau_{n-1}, \mu) + \delta \int_{\Lambda \times \mathscr{M}(\Lambda)} W_n^n(\tilde{\alpha}, \tau, \tilde{\mu}) \mathbf{P}(d(\tilde{\alpha}, \tilde{\mu}) | \alpha, a, \tau_{n-1}, \mu).$$

To see that $V_{n-1}^n(\alpha, a, \tau, \mu)$ is continuous in (α, a) and norm continuous in τ, observe that

$$\begin{aligned}
&\int sup_y | V_{n-1}^n(y, \tau^k, \mu) - V_{n-1}^n(y, \tau, \mu) | \psi(d\mu) \\
&\leq \int sup_y | u(\alpha, a, \tau_{n-1}^k, \mu) - u(\alpha, a, \tau_{n-1}, \mu) | \psi(d\mu) \\
&+ \delta \int sup_\alpha \left| \int W_n^n(\tilde{\alpha}, \tau^k, \tilde{\mu}) \mathbf{P}(d(\tilde{\alpha}, \tilde{\mu}) | \alpha, a, \tau_{n-1}^k, \mu) \right. \\
&\left. - \int W_n^n(\tilde{\alpha}, \tau, \tilde{\mu}) \mathbf{P}(d(\tilde{\alpha}, \tilde{\mu}) | \alpha, a, \tau_{n-1}, \mu) \right| \psi(d\mu).
\end{aligned}$$

The first term on the right goes to 0, by norm continuity of u. The second term is bounded from above by

$$\begin{aligned}
&\int sup_y \left| \int W_n^n(\tilde{\alpha}, \tau^k, \tilde{\mu}) \mathbf{P}(d(\tilde{\alpha}, \tilde{\mu}) | \alpha, a, \tau_{n-1}^k, \mu) \right. \\
&\left. - \int W_n^n(\tilde{\alpha}, \tau, \tilde{\mu}) \mathbf{P}(d(\tilde{\alpha}, \tilde{\mu}) | \alpha, a, \tau_{n-1}^k, \mu) \right| \psi(d\mu) \\
&+ \int sup_y \left| \int W_n^n(\tilde{\alpha}, \tau, \tilde{\mu}) \mathbf{P}(d(\tilde{\alpha}, \tilde{\mu}) | \alpha, a, \tau_{n-1}^k, \mu) \right. \\
&\left. - \int W_n^n(\tilde{\alpha}, \tau, \tilde{\mu}) \mathbf{P}(d(\tilde{\alpha}, \tilde{\mu}) | \alpha, a, \tau_{n-1}, \mu) \right| \psi(d\mu).
\end{aligned}$$

The second of these terms converges to 0, since $\mathbf{P}(d(\tilde{\alpha}, \tilde{\mu}) | \alpha, a, \tau_{n-1}, \mu)$ is norm continuous on measurable functions. The first of these two terms is bounded from above by

$$\begin{aligned}
&\int sup_y \int sup_{\tilde{\alpha}} | W_n^n(\tilde{\alpha}, \tau^k, \tilde{\mu}) - W_n^n(\tilde{\alpha}, \tau, \tilde{\mu}) | \mathbf{P}(d(\tilde{\alpha}, \tilde{\mu}) | y, \tau_{n-1}^k, \mu) \psi(d\mu) \\
&\leq \int_\mu \int_{\tilde{\mu}} sup_{\tilde{\alpha}} | W_n^n(\tilde{\alpha}, \tau^k, \tilde{\mu}) - W_n^n(\tilde{\alpha}, \tau, \tilde{\mu}) | b \psi(d\mu) \psi(d\tilde{\mu}).
\end{aligned}$$

The latter term converges to 0, hence $V_{n-1}^n(\alpha, a, \tau, \mu)$ is norm continuous in τ. Finally, continuity of $V_{n-1}^n(\alpha, a, \tau, \mu)$ in $y = (\alpha, a)$ follows directly from continuity of $u(\alpha, a, \tau_{n-1}, \mu)$ in y and since $\mathbf{P}(d(\tilde{\alpha}, \tilde{\mu}) | y, \tau_{n-1}^k, \mu)$ is assumed continuous in y on measurable functions.

Let $W_{n-1}^n(\alpha, \tau, \mu) = max_a V_{n-1}^n(\alpha, a, \tau, \mu)$, and proceed inductively to define sequences of functions, $\{V_t^n(\alpha, a, \tau, \mu)\}_{t=1}^n$ and $\{W_t^n(\alpha, \tau, \mu)\}_{t=1}^n$. As in the proof of theorem 3, the limits $\lim_n V_t^n(\alpha, a, \tau, \mu)$ and $\lim_n W_t^n(\alpha, a, \tau, \mu)$ exist and are norm continuous in τ and continuous in (α, a) and α respectively.

Next, observe that the intertemporal consistency conditions satisfy upper-hemicontinuity. To see this let $\tau^n \to \tau$ and let $\tilde{\tau}^n$ be a consistent sequence in the range of the correspondence with $\tilde{\tau}^n \to \tilde{\tau}$. Thus, considering period t, $\forall f \in C(\mathscr{M}(\Lambda))$, $\forall g \in \mathbf{L}_1(\mathscr{M}(\Lambda), \mathscr{R}, \psi)$

$$\int_{\mathscr{M}(\Lambda)} \tilde{\tau}_{t+1}(f, \Lambda; \mu) g(\mu) \psi(d\mu) = \int \mathbf{P}_{\mathscr{M}\xi}(f, \Lambda; y, \tau_t, \mu) g(\mu) \psi(d\mu), \forall y \in Y.$$

$\int_{\mathscr{M}(\Lambda)} \tilde{\tau}_{t+1}(f, \Lambda; \mu) g(\mu) \psi(d\mu)$ converges to $\int_{\mathscr{M}(\Lambda)} \tau_{t+1}(f, \Lambda; \mu) g(\mu) \psi(d\mu)$, in view of the topology on τ. Comparing $\int \mathbf{P}_{\mathscr{M}\xi}(f, \Lambda; y, \tau_t^n, \mu) g(\mu) \psi(d\mu)$ and $\int \mathbf{P}_{\mathscr{M}\xi}(f, \Lambda; y, \tau_t, \mu) g(\mu) \psi(d\mu)$, the difference (in absolute value) converges to 0, by norm continuity of $\mathbf{P}_{\mathscr{M}\xi}$.

Finally, we construct the best response mapping in exactly the same way as in theorem 3. A consistent strategy $\hat{\tau}$ is a best response (with consistent marginal λ on $\mathscr{M}(\Lambda)$: $\hat{\tau}(\bullet, \Lambda : \mu) = \lambda_t(\bullet : \mu)$, ψ *a.e.* μ),

$$\int_{\mathscr{M}(\Lambda)} \int_Y V_t(\alpha, a, \tau, \mu) \hat{\tau}_t(dy; \mu) \psi_t(d\mu) = \int_{\mathscr{M}(\Lambda)} \int_\Lambda W_t(\alpha, \tau, \mu) \lambda_t(d\alpha; \mu) \psi(d\mu).$$

The reasoning given in the proof of theorem 3, establishes existence here also. ■

References

Bergin, J.: A characterization of sequential equilibrium strategies in infinitely repeated incomplete information games. J. Econ. Theory **47**, 51–65 (1989)

Bergin, J., Bernhardt, D.: Business cycles, thin resale markets and Darwinian competiton. Mimeo, Queen's University, 1993

Bergin, J., Bernhardt, D.: Existence of equilibrium in anonymous sequential games with general state space of aggregate uncertainty. CORE discussion paper 9043, 1990

Bergin, J., Bernhardt, D.: Anonymous sequential games with aggregate uncertainty. J. Math. Econ. **21**, 543–562 (1992)

Davis, S., Haltiwanger, J., Schuh, S.: Job creation and destruction in U.S. manufacturing: 1972–1988. U.S. Census Bureau Monograph, 1993

Diestel, J., Uhl, J. J.: Vector measures. AMS Mathematical Surveys, No. 15, 1977

Duffie D., Geanakoplos, J., Mas-Colell, A., MacLennan, A.: Stationary Markov equilibria. Econometrica, forthcoming

Feldman, M., Gilles, C.: An expository note on individual risk without aggregate uncertainty. J. Econ. Theory **35**, 26–32 (1985)

Gouge, R., King, I.: Job reallocation, worker search and unemployment over the business cycle. Mimeo University of Victoria, 1994

Hopenhayn, H.: Entry, exit, and firm dynamics in long run equilibrium. Econometrica **60**, 1127–1150 (1992)

Jovanovic, B.: Selection and the evolution of industry. Econometrica **50**, 649–670 (1982)

Jovanovic, B., MacDonald, G. M.: Competitive diffusion. J. Pol. Econ. **102**, 24–52 (1994)

Jovanovic, B., Rosenthal, R. W.: Anonymous sequential games. J. Math. Econ. **17**, 77–88 (1988)

Mertens, J. F.: Repeated games. Proc. Int. Congr. Math. (1986)

Mertens, J. F., Parthasarathy, T.: Existence and characterization of Nash equilibria for discounted stochastic games. CORE DP 8750, 1988

Mertens, J. F.: Stochastic games. Mimeo 1989

MasColell, A.: On a Theorem of Schmeidler. J. Math. Econ. **13**, 201–206 (1984)