

EQUILIBRIA FOR DISCOUNTED STOCHASTIC GAMES

JEAN-FRANÇOIS MERTENS AND T.E.S. PARTHASARATHY
CORE, Université Catholique de Louvain
Louvain-la-Neuve, Belgium

Abstract. We prove the existence of subgame-perfect equilibria for discounted stochastic games with general state and action sets, using minimal assumptions (measurability as a function of states, and for each fixed state, compactness of action sets and continuity on those)—except for the rather strong assumption that the transition probabilities are norm-continuous functions of the actions.

1. Introduction

Equilibria for discounted stochastic games were not known to exist in the presence of uncountable state spaces. We prove this result, and get as an additional bonus that arbitrary compact action sets can also be allowed, and that one can furthermore get the equilibria to be subgame perfect.

Such a result should be an “archetype” result for a satisfactory existence theorem for rational expectations equilibria, either by borrowing and adapting the techniques of the proof, and/or in the same way as Nash’s [6] existence proof led to the existence of economic equilibria (e.g., [1]), by reinterpreting the economic equilibrium problem as a game, and/or by interpreting discounted stochastic games “à la Shapley–Shubik” as a general formulation of any strategic version of the rational expectations equilibrium problem, whose equilibria should converge [4] to the true rational expectations equilibria. Indeed, a central difficulty in establishing rational expectations equilibria—just like here—is to obtain that agents really use strategies, i.e., act only on the basis of their past physical observations, and not based on some mythical common expectation for the future that would suddenly become publicly available at that stage. (This would rather lead to “extensive form correlated equilibria” [3].) The existence of such a common expectation for the future will clearly be a consequence (by equi-

librium, symmetric information is the basic assumption), but it should not be a primitive datum.

It is this strategic type of interpretation we have chiefly in mind, both in the effort to obtain a reasonable set of assumptions and in the interpretation of the results.

It is true we only obtain subgame-perfect equilibria—there is no indication that it might be possible to have them in addition stationary. But the equilibria we obtain are stationary in the traditional sense of the rational expectations literature, in that they are stationary functions of the current state and the current common expectation about the future (the expected vector payoff from now on); cf the proof of Section 6, Lemma 11: $\tau(h, \omega) = \sigma(p, \omega)$, where $p = p(h, \omega)$ is the expected payoff vector for the future, ω is the current state, h the past history, and σ a fixed measurable function.

Before presenting the assumptions and the proof in more detail, we want to point out here that we use one strong assumption: that the distribution of tomorrow's state depends norm-continuously on today's actions. This will typically not be satisfied in the type of economic models just mentioned, if the future state is just some deterministic function of today's state and actions: some stochastic element has to be present, as it indeed usually is, e.g., in the form of a normal disturbance. The assumption is clearly vacuous for finite action sets.

The simplest model where, because of the failure of this assumption, we don't know whether equilibria exist—and which in all other respects is as well behaved as one can desire—is the following: the state space is the Cantor set $C = F^{\mathbb{Z}}$ (where F is any finite set—or group, or field if more structure than the shift operation is desired for specifying a counterexample); each player n 's strategy set S_n is the one-point compactification \mathbb{Z} of the integers; each player's payoff function today and the probability that tomorrow's state belongs to some Borel subset of C are continuous functions of today's state and actions, i.e., on $C \times \prod_n S_n$. All players discount future payoffs with the same positive discount factor. Note in particular that it follows in this model that all transitions are absolutely continuous with respect to a single measure μ (in such cases, important simplifications become possible in our proof)—and are even uniformly integrable in $L_1(\mu)$. Also, one has the best possible topological assumptions; and finally, if the expected payoff from tomorrow on is a given measurable function, its expectation today will be jointly continuous in the actions (and the state), so that today's game will satisfy all usual conditions for the existence of Nash equilibria.

This is at first sight all we need for our type of proof to work: indeed, the basic idea of the proof is to work by backward induction. Assume a

measurable, compact-valued correspondence N_0 from state space to payoff space—e.g., $N_0(\omega)$ is the set of all payoff vectors between $-A$ and $+A$, if A is the maximum in absolute value of all single-stage payoffs. Define then $N_{k+1}(\omega)$ inductively as the set of all Nash equilibrium payoffs at ω when the payoffs at tomorrow's state $\tilde{\omega}$ are some measurable selection from $N_k(\tilde{\omega})$. The inclusion of N_{k+1} in N_k follows by induction. This suggests defining the “candidate set” of equilibrium payoffs N_∞ as $\cap_k N_k$, at least with finite action sets. But with infinite action sets a closure operation is typically required here. Indeed, since the set of Nash payoffs of a game is not convex, $N_k(\tilde{\omega})$ has no reason to be convex-valued, so its set of selections is typically noncompact, for no relevant topology. Yet compactness of the $N_k(\omega)$ cannot be dispensed with, to insure $N_\infty(\omega) \neq \emptyset$. Thus strong (relative) compactness of the set of “today's games” is required. This in turn requires our continuity assumption on the transition probability. In terms of this strong topology, the Nash correspondence is still uppersemi-continuous (u.s.c.), hence the compactness of $N_{k+1}(\omega)$, after allowing for the limiting games. Adding those limiting games is harmless (they correspond to measurable selections from the convex hull of N_k), because for any given actions today, only one distribution of tomorrow's state is involved, so that Lyapunov's theorem can be used to find, as a function of today's actions, another measurable selection from the set of future payoffs $N_k(\tilde{\omega})$, with the correct expectation for those given actions today (and this is basically what causes the nonstationarity of the strategies).

The fact that today one can get any payoff in N_{k+1} in equilibrium provided that tomorrow one can do so for N_k goes to the limit and yields that today one can get any payoff in N_∞ as an equilibrium provided one can do so tomorrow. Any strategy that implements this idea forever is then finally shown to yield indeed the payoff expected (because of the discounting), and hence to be indeed an equilibrium. Clearly, one still has to show that the N_k 's will be measurable, and mainly that it is possible to reconstruct measurable strategies that will yield any point in the candidate set as an equilibrium. (In particular, one has to show that one can do the above application of Lyapunov in a measurable way as a function of today's state and action and of the payoff expected from the future.) But the above is indeed the basic idea of this paper, and the solution to those measure-theoretic problems is taken from [5].

The construction of the candidate set can be found in Section 5, while the construction of the corresponding strategies is done in Section 6. The basic “strong compactness” property is established in Section 3, Proposition 4.

The above “simplest” model would imply weak relative compactness (in the space of continuous functions on $\Pi_n S_n$) of the set of “today's games”

obtained in the above way. But this is of no help in obtaining compactness for the set of Nash equilibria, nor even for their set of payoffs $N_{k+1}(\omega)$ —indeed, not even in the one-person case. For example, consider $f_m(0) = \frac{1}{2}$, $f_m(i) = \frac{m}{i}$ for $i > m$, $f_m(i) = 0$ otherwise. The sequence of continuous functions f_m converges pointwise (“weakly”) to the continuous function f_∞ , yet the optimal strategy $(m+1)$ for f_m does not converge to the optimal strategy (0) for f_∞ , and $\max f_m = \frac{m}{m+1}$ does not converge to $\max f_\infty = \frac{1}{2}$ (and the limit, 1, of those maxima is not even achievable as the supremum of some function in the closed convex hull of our sequence).

It follows from the approach just sketched that, when for each state the corresponding action sets are finite, there are no continuity or other restrictions as a function of the actions, so the problem can be reformulated taking as larger state space the space of all finite histories. On this larger state space, all our assumptions are still satisfied, and the theorem then provides a full characterization of the set of all subgame-perfect equilibria.

In Section 2, we present the basic data of the model: state space, action sets, transition probabilities and single-stage payoff function. In a stochastic game, there is another game matrix for each state—in particular different action sets. With a continuum of states, however, one has to express that those depend in a measurable way on the state. Hence we need a common embedding space, and to express the action sets as a measurable, compact-valued correspondence from the state space to this embedding space. Measurability of the graph is the most standard and the weakest assumption in such a setting. The theorem becomes both more precise and simpler, however, under a slightly stronger assumption that this correspondence is measurable as a map to the space of compact subsets of this embedding space. Section 2 elucidates the relationships between those two assumptions (with the related measurability assumptions on payoff functions and transition probabilities), shows how the result for the weaker assumptions will follow from that for the stronger assumptions, and goes a long way towards proving that it is sufficient to consider state spaces which are *separable* measurable spaces. This last point is finished in Section 4.

Finally, Section 3 deals with the assumptions on utilities. While we remain basically in a framework of (timewise) additively separable utilities, assumptions that further guarantee that each stage’s utility remain the same (uniformly bounded) function of current state and actions, and that all players at all times use the same discount factor, would be extremely restrictive and unpalatable. As soon as one wants to accommodate discount factors that vary, e.g., from agent to agent (or from date to date), it is more convenient to reformulate the problem over space-time as new state space, and to incorporate the effect of all discount factors into the single-stage payoff function itself. Such an operation always preserves all other

assumptions of the model, and a similar trick could be used to allow for state-dependent discount factors, or even single-stage payoff functions that depend also on the past history of states (not actions, however, if action sets are infinite).

We are thus led to consider, for each player, one single payoff function, depending only on current state and actions, and where his payoff for the whole game is thought of as the sum of the payoffs he receives each day. Under standard assumptions (uniformly bounded utilities each day, and constant positive discount factors), those sums will converge for any possible history, and the expectation of the sum will, for any probability distribution, be the sum of the expectations. But such assumptions are extremely hard to justify. For example, if one has indeed replaced the state space by the space of all finite state histories, a uniform boundedness assumption on utilities is completely meaningless, since it is destroyed when multiplying utilities by a positive function of the initial state, and this leaves all relevant decision problems unchanged. Also, an assumption that the sum of the payoffs converges for any possible history will be very restrictive for actual economic models: even if an agent has each day the same logarithmic utility function for current consumption, and a constant positive discount rate is used, there will always be some histories, however unlikely, where, due to extremely favorable outcomes of the previously mentioned normal disturbances, consumption will grow so fast that the sum of utilities will not converge.

We are thus led to define as payoff function for the stochastic game the sum of the expected payoffs corresponding to each stage. We assume those expectations to exist and those sums to converge for any strategy choices, and we even slightly strengthen those assumptions, so as to make sure that in addition the payoff to any mixed strategy is the corresponding expectation of payoffs to pure strategies. Such assumptions guarantee that all payoffs we consider will have the intended meaning, and be free from the above-mentioned drawbacks—both conceptual and in practical applications. A further assumption expresses that the payoff depends only a little on the far tail of the strategies, thus expressing an obvious implication of any discounting setup.

Finally, an effort is made to express the assumptions as much as possible in terms of pure strategies—and to deduce corresponding properties over mixed strategies; this is a general endeavor in game theory, since pure strategies are viewed as basic in the model, while mixed strategies are in some sense a fiction of the mind for obtaining a solution.

A first part of Section 3 is devoted to obtaining the basic consequences of those assumptions: that payoffs are well defined, for any strategy vector of whatever form; that it is sufficient to consider behavioral strategies, and

that the payoffs then satisfy a recursive relation; and finally, that even if one restricts oneself, as we do, to behavioral strategies satisfying the strictest measurability requirements, for the sake of the strength of the theorem, Nash equilibria in this restricted class of strategies would remain such if the class of strategies were relaxed in any meaningful sense. Thus both payoffs and Nash equilibria are completely unambiguous, and the recursive relation is established.

Several of the assumptions were needed only for the above purpose; for the rest of the proof only the recursive equation will be needed. So at this stage a new set of assumptions is introduced, much less restrictive than the previous set, but also less primitive in some sense, bearing directly on this recursive equation (Section 3, Proposition 3). Those are the only assumptions necessary for the sequel. They allow in particular some asymptotic part in the payoff, and are such that, in the new formulation, an extension to general payoff functions, not additively separable, becomes a matter of immediate generalization (just reformulate the recursion formulas so as to include, in the expected payoff after every finite history, not only the future expected payoff, but also the payoff accumulated in the past).

The last part of Section 3 is then devoted to establishing, under those new assumptions, the basic properties—strong compactness in Proposition 4, recursive equation and continuity at ∞ in Proposition 5; then to showing that the concept of subgame-perfect equilibrium is also completely unambiguous, and that it coincides with the backward induction equilibria.

At this stage one is ready to finish the separability issue in Section 4, and then to move to the previously sketched core of the proof in Sections 5 and 6.

However, because of our very general assumptions on the payoff function (even without the reformulation), the above-sketched construction of $N_0(\infty)$ is no longer adequate. To take a trivial example, an expected payoff identically equal to 1 (which is paid at “the end of the game”), together with a payoff for each stage identically zero, satisfies all assumptions of Proposition 3 in Section 3. Yet if N_∞ is going to be the right thing, one needs $N_0(\omega)$ to be just the payoff vector 1, with nothing else. Thus in Section 5 we first have to construct explicitly a measurable compact-valued correspondence $N_0(\omega)$, such that any payoff in $N_0(\omega)$ is feasible starting from ω , and which is sufficiently large to ensure that one will indeed have $N_1(\omega) \subseteq N_0(\omega)$. This $N_0(\omega)$ is constructed by a similar induction to the one previously sketched for $N_\infty(\omega)$, only this time because it is an increasing sequence, one looks at all feasible payoffs and not only at the equilibrium payoffs, and one starts the induction with the actual payoff to some stationary strategy.

Those are the basic ideas in the paper. The real thing follows.

2. The Model

2.1. THE DATA

- (a) The state space is a measurable space (Ω, \mathcal{A}) .
- (b) For each player n , his strategy space is a measurable map S_n from (Ω, \mathcal{A}) to $\mathcal{K}_{\bar{S}_n}^*$, i.e., the space of nonempty compact subsets of \bar{S}_n endowed with the Hausdorff topology (and the corresponding Borel σ -field), where \bar{S}_n is a separable metric space (or any Souslin space—a regular Hausdorff space which is the continuous image of a Borel subspace of a compact metric space) or at least:
- (b') $S_n(\omega) \in \mathcal{K}_{\bar{S}_n}^*$, \bar{S}_n is a Lusin space (like a Souslin space, but the continuous mapping is required to be one to one) with Borel σ -field \mathcal{S}_n , and the graph of S_n is $\mathcal{A} \otimes \mathcal{S}_n$ -bianalytic ([5], Appendix, §1).
- (c) Define $(\bar{S}, \mathcal{S}) = \bigotimes_n (\bar{S}_n, \mathcal{S}_n)$, and $S(\omega) = \Pi_n S_n(\omega)$ and let $(G, \mathcal{G}) \subseteq (\bar{S} \times \bar{S}, \mathcal{A} \otimes \mathcal{S})$ denote the graph of S : $G = \{(\omega, s) \mid s \in S(\omega)\}$.
- (d) A transition probability from (G, \mathcal{G}) to (Ω, \mathcal{A}) , $p(A \mid g)$ is given; i.e., $\forall g \in G$, $p(\cdot \mid g)$ is a probability distribution on (Ω, \mathcal{A}) and $\forall A \in \mathcal{A}$, $p(A \mid g)$ is \mathcal{G} -measurable.
- (e) For each player n , a measurable payoff function u_n on (G, \mathcal{G}) is given or at least:
- (d'), (e') The measurability requirements on $p(A \mid g)$ and on $u_n(g)$ can be weakened to having an analytic graph in $(G, \mathcal{G}) \times (\mathbb{R}, \text{Borel sets})$.
- (f) For each ω , the functions $u_n(\omega, s)$ and $p(\cdot \mid \omega, s)$ are continuous on $S(\omega)$, using the norm topology for measures on the state space.

2.2. THE GAME

An initial state is chosen according to some given probability distribution μ on (Ω, \mathcal{A}) .

At each stage, each of the players is first informed of the current state $\omega \in \Omega$; next they all simultaneously choose an action— $s_n \in S_n(\omega)$ for player n ; given the point $g \in G$ thus obtained, each player n receives his payoff $u_n(g)$, and is informed of $s = (s_1, s_2, s_3, \dots)$; next a new state is selected according to $p(\cdot \mid g)$, and the game moves to the next state.

The players use behavioral strategies; cf Section 3 for more details.

For any strategy vector, each player n can compute his expected payoff u_n^t relative to each stage t . His overall payoff is $\sum_t u_n^t$, which we will assume to be well defined.

2.3. THE RELATION BETWEEN ASSUMPTIONS (A) AND (A')

The main reason for mentioning the possibility of the weaker requirements (A') (i.e., (a), (b'), (c), (d'), (e') and (f)) instead of the stronger requirements (A) (i.e., (a), (b), (c), (d), (e) and (f)) is that in particular (b) is unnaturally strong, in that it depends on the topology of the embedding space \bar{S}_n , instead of just its Borel structure and the topology of the sets $S_n(\omega)$, and that an assumption of measurability of the graph is more classical. Typical examples of nonmetrizable Lusin spaces that could plausibly occur as embedding spaces are separable Banach spaces or their duals in their weak (resp. weak*) topology.

The next proposition relates the two sets of requirements. We will not distinguish between two models that are identical except for the embedding spaces \bar{S}_n , as long as the topology of the sets $S_n(\omega)$ and the Borel structure on $\cup_\omega S_n(\omega)$ remain the same, since the sets of strategies, and hence the payoff functions, will depend only on this Borel structure.

Proposition 1 *a) There is no loss of generality in assuming $\bar{S}_n = [0, 1]^\infty$ — hence compact metric.*

b) (A) implies (A').

c) Under (A'), there exists a minimal σ -field \mathcal{F} on Ω such that the Assumptions (A) are satisfied when \mathcal{F} replaces \mathcal{A} . Furthermore, \mathcal{F} is separable.

d) Under (A'), (A) is satisfied when the bianalytic σ -field $\mathcal{B} = \mathcal{B}(\mathcal{A})$ ([5], Appendix, §1) replaces \mathcal{A} .

e) In particular, the separable σ -field \mathcal{F} satisfies $\mathcal{F} \subseteq \mathcal{A}$ under (A), and $\mathcal{F} \subseteq \mathcal{B}$ under (A').

f) $\mathcal{B} = \mathcal{A}$ when (Ω, \mathcal{A}) is a Blackwell space—and (Ω, \mathcal{F}) is then a Blackwell space too.

((Ω, \mathcal{A}) is a Blackwell space if it is (measure-theoretically) isomorphic—after identification of \mathcal{A} -equivalent points of Ω —to an analytic subset ([5], Appendix, §1) of a Lusin space, with the Borel σ -field. Or equivalently, if \mathcal{A} is separable, and any real-valued measurable function has an analytic range.)

Proof. (a) If \bar{S}_n is separable metric, it can be embedded as a subspace of $[0, 1]^\infty$; use [5], Proposition 6.g to conclude in the case of requirement (b). Assume thus \bar{S}_n Souslin or Lusin. By [5], Proposition 4 there exists a weaker, separable metric topology on \bar{S}_n . Compact subsets of \bar{S}_n are still compact in the new topology, and the two topologies have the same Borel sets (by the first separation theorem for analytic sets). Thus all requirements that were valid are still valid under the new topology; for requirement (b) use, e.g., [5], Proposition 9.a. Thus we can assume \bar{S}_n in addition separable and metric. In the case of requirement (b), this finishes the proof. For

(b'), note that it is then a subspace of $[0, 1]^\infty$, and, being Lusin, a Borel subset, so that requirement (b')—and clearly the other requirements—remains valid if we think of \bar{S}_n as $[0, 1]^\infty$. This proves (a).

(b) Fix for each k a finite open covering of $\bar{S}_n = [0, 1]^\infty$ by balls of radius k^{-1} , and $\forall K \in \mathcal{K}_{\bar{S}_n}^*$; let $\varphi_k(K)$ denote the union of the closures of those balls that intersect K . Since the set of K 's that intersect a fixed ball is open, φ_k is a Borel map from $\mathcal{K}_{\bar{S}_n}^*$ to itself, and φ_k converges to the identity. If $S_n(\omega)$ is a measurable map, so is $(\varphi_k \circ S_n)$, and being a step function, the latter clearly has a measurable graph. Thus the graph of S_n , being the intersection of those graphs, is also measurable, and a fortiori bianalytic: (b) \Rightarrow (b'). By the same argument, (d) \Rightarrow (d') and (e) \Rightarrow (e').

(c) $\mathcal{K}_{[0,1]^\infty}^*$ is compact metric, so the σ -field \mathcal{A}_0 that makes all maps S_n measurable is separable. \mathcal{A}_0 depends only on the topology of the sets $S_n(\omega)$ and on the Borel structure on $\cup_\omega S_n(\omega)$. Indeed, by [5], Proposition 7.b, there exists a sequence of \mathcal{A}_0 -measurable selections from $S_n(\omega)$, which are for each ω dense in $S_n(\omega)$. (If \bar{S}_n is a nonmetrizable Souslin space, go to a weaker metrizable topology ([5], Proposition 4) before using Proposition 3.b; this preserves the topology of the $S_n(\omega)$ and the Borel structure on $\cup_\omega S_n(\omega)$.) The properties that we mentioned of the sequence depend only on the topology of the sets $S_n(\omega)$ and on the Borel structure on $\cup_\omega S_n(\omega)$. Conversely, as soon as there is such a sequence, [5], Corollary 8bis, (a) implies that the map S_n is measurable. (If \bar{S}_n were a nonmetrizable Souslin space, we would conclude that S_n is Borel measurable when \bar{S}_n is endowed with any weaker metrizable topology; since ([5], Proposition 4) any open set in \bar{S}_n is open in some such topology, we would still conclude that S_n is measurable when $\mathcal{K}_{\bar{S}_n}$ is endowed with the Effrös σ -field (using [5], Proposition 4, this σ -field can be equivalently described as the σ -field generated by the sets $\{K \mid K \subseteq O\}$ or by the sets $\{K \mid K \cap O \neq \emptyset\}$, O being an open set, or still equivalently, a closed set). Apparently, in that case our assumption may still be somewhat stronger than needed, but no matter—it is obvious at this stage anyway that in fact we only used the Effrös-measurability of the maps $S_n(\omega)$; hence requirement (b) could be weakened accordingly.) In particular, G is $\mathcal{A}_0 \times \mathcal{S}$ -measurable in $\Omega \times \bar{S}$ ([5], Proposition 9.b).

Note that $\mathcal{A}_0 \subseteq \mathcal{A}$ under Assumptions (A), and $\mathcal{A}_0 \subseteq \mathcal{B}(\mathcal{A})$ (by [5], Proposition 1) under Assumptions (A').

Now consider a real-valued function $f(\omega, s)$ on G , which is continuous in $s \in S(\omega)$ for each ω , and has an analytic graph. Let $s_i(\omega)$ denote the above-mentioned sequence of \mathcal{A}_0 -measurable pure strategy vectors which is, for each ω , dense in $S(\omega)$. Then the graph $\{(x, \omega, s) \mid s = s_i(\omega), x = f(\omega, s)\}$ is analytic, as an intersection of two such sets. Being the graph of a function, this function is bianalytic—i.e., $\mathcal{B}(\mathcal{A})$ -measurable—since the inverse image

of each Borel set is analytic, by projection. Note that, if f is \mathcal{G} -measurable and $\mathcal{A}_0 \subseteq \mathcal{A}$, this function will be \mathcal{A} -measurable, by composition. The continuity of f in s implies, as before, by [5], Corollary 8bis,a, that the map $\omega \rightarrow \gamma_f(\omega) = \{(x, s) \mid s \in S(\omega), x = f(\omega, s)\}$ is a $\mathcal{B}(\mathcal{A})$ -measurable (\mathcal{A} -measurable if f is G -measurable and $\mathcal{A}_0 \subset \mathcal{A}$) map from Ω to $\mathcal{K}_{\bar{S} \times \mathbb{R}}^*$. Denote by \mathcal{A}_f the minimal σ -field for which this measurability is true. Since $\mathcal{K}_{\bar{S} \times \mathbb{R}}^*$ is locally compact with a countable basis, \mathcal{A}_f is separable. Our proof also shows that $\mathcal{A}_f \subseteq \mathcal{B}(\mathcal{A})$, and $\mathcal{A}_f \subseteq \mathcal{A}$ under Assumptions (A) when f is \mathcal{G} -measurable; and that \mathcal{A}_f is included in the σ -field generated by the maps $(s_i(\omega), f(\omega, s_i(\omega)))$. Further, since $S(\omega)$ is the projection of $\gamma_f(\omega)$, $S(\omega)$ is \mathcal{A}_f -measurable ([5], Proposition 6.5), hence $\mathcal{A}_0 \subseteq \mathcal{A}_f$. Thus $s_i(\omega)$ is \mathcal{A}_f -measurable, and similarly $f(\omega, s_i(\omega))$ is \mathcal{A}_f -measurable, as the composition of $\omega \rightarrow \gamma_f(\omega) \cap (\mathbb{R} \times \{s_i(\omega)\})$, which is \mathcal{A}_f -measurable by [5], Proposition 6.e since both $\gamma_f(\omega)$ and $s_i(\omega)$ are so, with the projection to the first factor space \mathbb{R} . Thus \mathcal{A}_f is exactly the σ -field generated by the maps $(s_i(\omega), f(\omega, s_i(\omega)))$, whatever may be the sequence of \mathcal{A}_0 -measurable pure strategy vectors $s_i(\omega)$ such that, for each ω , the sequence $s_i(\omega)$ is dense in $S(\omega)$. As before in the definition of \mathcal{A}_0 , this characterizes \mathcal{A}_f in a way that depends only on the topology of the sets $S(\omega)$ and on the measurable structure on $\cup_\omega S(\omega)$.

Now let $A_0^n = \emptyset$, and for $1 \leq i \leq n$, $A_i^n = \{(\omega, s) \mid (\omega, s) \notin \cup_{j < i} A_j^n \text{ and } d(s, s_i(\omega)) \leq d(s, s_j(\omega)) \forall j > i, j \leq n\}$. The A_i^n ($i = 1, \dots, n$) clearly form an $\mathcal{A}_0 \otimes \mathcal{S}$ -measurable partition of $\Omega \times \bar{S}$. Let $f_n(\omega, s) = f_n(\omega, s_i(\omega))$ for $(\omega, s) \in A_i^n$: $f_n(\omega, s)$ is then clearly an $\mathcal{A}_f \otimes \mathcal{S}$ -measurable function on $\Omega \times \bar{S}$. Let $\varphi(\omega, s) = \liminf_{n \rightarrow \infty} f_n(\omega, s)$: φ is also $\mathcal{A}_f \otimes \mathcal{S}$ -measurable, and the continuity of f in s implies that, for $s \in S(\omega)$, $\varphi(\omega, s) = f(\omega, s)$. Since G itself is $\mathcal{A}_f \otimes \mathcal{S}$ -measurable, it indeed follows that f is $\mathcal{A}_f \otimes \mathcal{S}$ -measurable on G . Finally, since we have seen above that $\mathcal{A}_f \subseteq \mathcal{A}$ as soon as f is $\mathcal{A} \otimes \mathcal{S}$ -measurable on G and G itself is $\mathcal{A} \otimes \mathcal{S}$ -measurable, it follows that \mathcal{A}_f is also the minimal σ -field containing \mathcal{A}_0 such that f is $\mathcal{A}_f \otimes \mathcal{S}$ -measurable on G .

Now let \mathcal{A}_1 denote the separable σ -field generated by the σ -fields \mathcal{A}_{u_n} for all players n . Note now that, for $B \in \mathcal{A}$, $p(B \mid g)$ satisfies our assumptions on the function f ; it is therefore in particular $\mathcal{B}(\mathcal{G})$ -measurable. Thus p is a transition probability from $(G, \mathcal{B}(\mathcal{G}))$ to (Ω, \mathcal{A}) . By [5], Proposition 3.a.1 it is a transition probability from $(G, \mathcal{B}(\mathcal{G}))$ to (Ω, \mathcal{B}) . Thus, for $B \in \mathcal{B}$, $p(B \mid g)$ is $\mathcal{B}(\mathcal{G})$ -measurable—in particular, has an analytic graph—and is continuous in s for fixed ω (because only countably many values of s are involved in the proof, and bianalytic sets are universally measurable). Thus $p(B \mid g)$ satisfies our assumptions for all $B \in \mathcal{B}$. If \mathcal{H} denotes a separable sub- σ -field of \mathcal{B} , denote by \mathcal{H}_0 a countable algebra that generates \mathcal{H} . Denote by $\mathcal{A}_{\mathcal{H}_0}$ the separable σ -field generated by all the \mathcal{A}_f for $f(g) = p(B \mid g)$

when B varies through \mathcal{H}_0 . A monotone class argument shows immediately that $p(B \mid g)$ is $\mathcal{A}_{\mathcal{H}_0} \times \mathcal{S}$ -measurable for all $B \in \mathcal{H}$; hence the separable σ -field $\mathcal{A}_{\mathcal{H}_0}$ depends only on \mathcal{H} : the minimal σ -field $\mathcal{A}_{\mathcal{H}}$ such that $p(B \mid g)$ is $\mathcal{A}_{\mathcal{H}} \otimes \mathcal{S}$ -measurable for all $B \in \mathcal{H}$ exists and is a separable sub- σ -field of \mathcal{B} .

Define thus \mathcal{A}_{n+1} inductively as the σ -field spanned by \mathcal{A}_1 and $\mathcal{A}_{\mathcal{A}_n}$: the σ -field \mathcal{F} spanned by the union of this increasing sequence of σ -fields is a separable σ -field containing \mathcal{A}_1 , and such that $p(B \mid g)$ is $\mathcal{F} \otimes \mathcal{S}$ -measurable for any B in the algebra $\cup_n \mathcal{A}_n$, hence, by a monotone class argument, for any $B \in \mathcal{F}$.

This proves (c), and at the same time (d) (and hence (e)). As for (f), its first part follows from the first separation theorem for analytic sets, and the second part then follows from the separability of \mathcal{F} (preferably using the second of the two equivalent definitions).

Remark 1. The definitions of strategies will be such that each player's strategy set varies monotonically with the σ -field \mathcal{A} on Ω . We want to show the existence of subgame-perfect equilibria—where the measurability assumptions are w.r.t. \mathcal{A} under requirements (A), w.r.t. \mathcal{B} under (A'). It will therefore be sufficient to show existence of a vector of strategies for \mathcal{F} , which are still a subgame-perfect equilibrium for \mathcal{B} (because the set of subgames too will vary—if at all—monotonically with the σ -field \mathcal{A}). Since the definition of \mathcal{F} itself is not affected either when \mathcal{A} is replaced by \mathcal{B} , and since the requirements (A) are satisfied when \mathcal{A} is replaced by either \mathcal{F} or \mathcal{B} , it follows that henceforth we can—and will—assume that the requirements (A) are satisfied. One will only have to remember that, in the case of (A'), \mathcal{A} in fact stands for the original $\mathcal{B}(\mathcal{A})$.

Remark 2. Thus, even for completely pathological state spaces (Ω, \mathcal{A}) , and just under the Assumptions (A'), we obtain subgame-perfect equilibria satisfying very stringent measurability requirements: it is just the σ -field \mathcal{A} that has to be extended—and not, e.g., the product σ -fields on finite histories—and it has to be extended only by bianalytic sets: this is the most conservative extension of a σ -field; it coincides in all classical cases with the original σ -field, and even in the very pathological cases it is just the right measure-theoretic analog of the effectively computable sets.

Remark 3. In the course of the above proof, we have also shown the following.

Corollary 1 *If f is a measurable function from (G, \mathcal{G}) to a separable metric space X , which is continuous in $s \in S(\omega)$ for each ω , then the map $\gamma_f : \omega \rightarrow \{(x, s) \mid s \in S(\omega), x = f(\omega, s)\}$ is measurable from (Ω, \mathcal{A}) to $\mathcal{K}_{X \times \bar{S}}^*$.*

Corollary 2 *There exists a sequence of pure strategy vectors s^i (measurable selections from $S(\omega)$) such that, for each ω , the sequence $s^i(\omega)$ is dense in $S(\omega)$.*

Remark 4. Observe that the proof in fact exhibited a countable algebra \mathcal{F}_0 , which generates \mathcal{F} , and which could be thought of as the basis of clopen sets for some separable, pseudo-metrizable topology on Ω (i.e., it is an embedding into the Cantor set), for which the map $\omega \rightarrow S(\omega)$ is continuous, as are the functions $u(g)$ and $p(f \mid g)$ (using the product topology on G) for all (\mathcal{F}_0) -continuous bounded functions f . Many variants of such a topological construction are possible; they could even be used to embed (Ω, \mathcal{F}) into a Polish space $\tilde{\Omega}$, and extend payoff function and transition probabilities to $\tilde{\Omega}$, such that $\tilde{u}(\cdot)$ and $\tilde{p}(f \mid \cdot)$ become continuous on \tilde{G} , for any bounded continuous f on $\tilde{\Omega}$ (and $S(\cdot)$ is continuous on $\tilde{\Omega}$). But, except in the case where the measures $p(\cdot \mid g)$ are dominated, we cannot guarantee that $\tilde{p}(\cdot \mid \tilde{\omega}, s)$ is still norm-continuous. This is apparently why, in the undominated case, allowing for general measurable spaces (Ω, \mathcal{A}) is really more general than just allowing for Blackwell spaces or standard Borel spaces. It is also why in the sequel we will have no use for such topological constructs.

2.4. A MORE INTRINSIC REFORMULATION: ASSUMPTIONS (\tilde{A}) AND (\tilde{A}')

The formulation of requirements (b) and (b') in the data of the model is not optimal. The following is at the same time easier, more general, and more intrinsic (in that it needs only the Borel structure on the embedding space of the action sets, and not its topology).

We need for each player n a measurable embedding space $(\bar{S}_n, \mathcal{S}_n)$, where the σ -field \mathcal{S}_n is both separable (countably generated) and separating (for each pair of distinct points there exists a measurable set containing one and not the other).

We also need a map $S_n(\omega)$ which assigns to each state ω a pair formed of a nonempty subset of \bar{S}_n , together with some compact topology on this subset.

As a link between the topologies and the measurable structures we require that the σ -field \mathcal{S}_n be generated by the measurable real-valued functions on $(\bar{S}_n, \mathcal{S}_n)$ with a continuous restriction to each compact set $S_n(\omega)$.

Lemma 1 *a) The above conditions are necessary and sufficient for the existence of a one-to-one map φ_n from \bar{S}_n into the unit cube $[0, 1]^{\mathbb{N}}$ (endowed with the usual topology and the Borel sets), which is an isomorphism of measurable structures and whose restriction to each compact set $S_n(\omega)$ is a homeomorphism.*

b) $\varphi_n(\bar{S}_n)$ is a Borel set if and only if $(\bar{S}_n, \mathcal{S}_n)$ is a standard Borel space (or finite, or countable).

Proof. Part a) is clear, if $(\bar{S}_n, \mathcal{S}_n)$ is a subset of the unit cube with the Borel sets, and if the sets $S_n(\omega)$ are compact for the subspace topology. Indeed, then \mathcal{S}_n is both separable and separating, and—since the continuous functions on the cube already generated the Borel σ -field—generated by the measurable functions with a continuous restriction to each set $S_n(\omega)$. Conversely, let A_k be a sequence of measurable sets that generate \mathcal{S}_n . Let \mathcal{F} denote the set of bounded real-valued measurable functions on $(\bar{S}_n, \mathcal{S}_n)$ with a continuous restriction to each $S_n(\omega)$ — \mathcal{S}_n is already generated by those. Thus, for each k , there exists a countable subset \mathcal{F}_k of \mathcal{F} such that A_k already belongs to the σ -field $\sigma(\mathcal{F}_k)$ generated by \mathcal{F}_k (because $\mathcal{S}_n = \sigma(\mathcal{F}) = \cup\{\sigma(\mathcal{D}) \mid \mathcal{D} \subseteq \mathcal{F}, \mathcal{D} \text{ countable}\}$ —the last equality because any countable union of countable sets is countable). Let f_i be an enumeration of $\cup_k \mathcal{F}_k$; we have $\mathcal{S}_n = \sigma(\{f_i \mid i \in \mathbb{N}\})$. There is no loss in scaling the f_i 's such that $0 \leq f_i(s) \leq 1$ for all s . Then let $\varphi_n(s) = (f_i(s))_{i \in \mathbb{N}} : \varphi_n$ maps \bar{S}_n into the unit cube. φ_n is measurable and its restriction to each set $S_n(\omega)$ is continuous because their composition with each coordinate projection is so. Consider thus the sub σ -field $\varphi_n^{-1}(\mathcal{B})$ of \mathcal{S}_n —where \mathcal{B} denotes the Borel σ -field on the cube. All functions f_i are measurable for this sub σ -field, hence $\mathcal{S}_n = \sigma(\{f_i \mid i \in \mathbb{N}\}) \subset \varphi_n^{-1}(\mathcal{B})$, thus $\mathcal{S}_n = \varphi_n^{-1}(\mathcal{B}) : \varphi_n$ is an isomorphism of measurable structures. Since \mathcal{S}_n separates points, it follows that φ_n is one to one. Thus the restriction of φ_n to each compact set $S_n(\omega)$ is a continuous one-to-one map to a Hausdorff topological space, hence is a homeomorphism with its image. This proves (a). Since φ_n is an isomorphism of measurable structures, (b) therefore amounts to showing that a subset of the unit cube is a Borel set if and only if it is either finite or countable or a standard Borel space. This is well known.

Up to now, the states ω serve just as a convenient way to index a family of subsets of \bar{S}_n , each endowed with some compact topology. The role of Assumptions (b) or (b') is to express the measurability of this. They become

- (\tilde{b}) $\{\omega \mid S_n(\omega) \cap U \neq \emptyset\}$ is measurable for each $U \in \mathcal{S}_n$ whose trace on each set $S_n(\omega)$ is open; and
- (\tilde{b}') $(\bar{S}_n, \mathcal{S}_n)$ is standard Borel (or finite, or countable), and the graph $G_n = \{(\omega, s) \mid s \in S_n(\omega)\}$ of S_n is bianalytic in $(\Omega \times \bar{S}_n, \mathcal{A} \otimes \mathcal{S}_n)$.

We finally relate those assumptions to the previous ones.

Proposition 2 a) Assumptions (\tilde{b}) and (\tilde{b}') are satisfied whenever (some version of) the corresponding Assumptions (b) and (b') are satisfied.

b) Under any identification φ_n as in Lemma 1 between \bar{S}_n and a subset of the unit cube, (\tilde{b}) is equivalent to (b), i.e., to the requirement that S_n be a measurable map to $\mathcal{K}_{\bar{S}_n}^*$ (endowed with the Hausdorff topology and

the corresponding Borel sets); and (\tilde{b}') is equivalent to (b') , i.e., to the requirement that G_n is bianalytic in the product of Ω with the unit cube (or with some Lusin subspace of it like \tilde{S}_n).

Proof. a) Under any set of assumptions, we ended up (Proposition 1.a) with a Borel isomorphism of \tilde{S}_n with a subset of the cube, which was a homeomorphism on each set $S_n(\omega)$. Thus the conditions of the present framework are satisfied.

Whenever some version of (b) was satisfied, we had the property that there exists a sequence of measurable selections $s^i(\omega)$ from $S_n(\omega)$, such that for each ω the sequence of values $s^i(\omega)$ is dense in $S_n(\omega)$ (cf proof of Proposition 1.c). But then $\{\omega \mid S_n(\omega) \cap U \neq \emptyset\} = \cup_i \{\omega \mid s^i(\omega) \in U\}$ for each set U which has an open trace on every set $S_n(\omega)$; hence this set is measurable as soon as U is.

Under (b') , we required \tilde{S}_n to be a Lusin space: this implies that it is a standard Borel space (or finite, or countable). The bianalyticity requirement is unchanged.

b) It remains thus to show that $\tilde{b} \Rightarrow b, \tilde{b}' \Rightarrow b'$. By Lemma 1.a, such identifications φ_n always exist. Condition (\tilde{b}) remains unchanged when the set \tilde{S}_n is increased or decreased, as it depends only on the trace of the σ -field on $\cup_\omega S_n(\omega)$. Hence, for (\tilde{b}) , we can assume \tilde{S}_n is the cube. Now consider a closed set C , and let $U_i = \{s \mid d(s, C) < 1/i\}$: we have $\{\omega \mid S_n(\omega) \cap C = \emptyset\} = \cup_i \{\omega \mid S_n(\omega) \cap U_i = \emptyset\}$, because $S_n(\omega) \cap C = \emptyset \Rightarrow d(S_n(\omega), C) > 0$. Thus $\{\omega \mid S_n(\omega) \cap C = \emptyset\}$ is measurable, i.e., writing U for the complement of C , we also get that, for any open set U , $\{\omega \mid S_n(\omega) \subseteq U\}$ is measurable. Now the Hausdorff topology on compact subsets of the cube is metrizable and separable; hence, to prove Borel measurability of a map to this space, it is sufficient to prove that the inverse images of some collection of basic open sets are measurable. Take those basic open sets of the form $\{S \mid S \subseteq U, S \cap U_i \neq \emptyset \forall i = 1, \dots, n\}$, where U and the U_i are open. Their inverse images are then $\{\omega \mid S_n(\omega) \subseteq U\} \cap [\cap_{i=1}^n \{\omega \mid S_n(\omega) \cap U_i \neq \emptyset\}]$, which is a finite intersection of sets of which we have proved the measurability. Hence (\tilde{b}) indeed implies (b) . That (\tilde{b}') implies (b') stems from part (b) of Lemma 1, since a Borel subset of the cube is a Lusin space, and from the fact that, if (B, \mathcal{B}) is a measurable space, with two subsets $B_1 \subseteq B_2 \in \mathcal{B}$, then B_1 is bianalytic in B if and only if it is so in B_2 —which stems in turn from the obvious analogous property for analytic sets, and from the inclusion of \mathcal{B} in the analytic sets.

It follows that we can use Assumptions (\tilde{b}) and (\tilde{b}') instead of (b) and (b') : the resulting set of assumptions requires no topology on the embedding spaces \tilde{S}_n , just their measurable structure—and the topology of the sets $S_n(\omega)$.

Note that the measurability requirement (\tilde{b}) has an unusual form: given a set S with a separable and separating σ -field \mathcal{S} , define a compatible family of compact subsets as a collection of pairs formed by a subset of S and a compact topology on this subset, such that the real-valued measurable functions on S with a continuous restriction to each of those compact subsets generate \mathcal{S} . Define the space \mathcal{K} of compact subsets as the set of all compatible families consisting of a singleton (i.e., they are the pairs formed by an element of \mathcal{S} and a compact metric topology on it whose Borel σ -field is the trace of \mathcal{S}). Compatible families are subsets F of \mathcal{K} , and have a corresponding σ -field \mathcal{T}_F generated by $\{X \in F \mid X \cap U \neq \emptyset\}$, where $U \in \mathcal{S}$ is such that $X \cap U$ is open in X for all $X \in F$.

Assumption (\tilde{b}) is that the map S_n be \mathcal{T}_F -measurable, where $F = \{S_n(\omega) \mid \omega \in \Omega\}$: the σ -field depends on the range of the map itself. However, it is clear that if F is a compatible family, and $G \subset F$, then G is also a compatible family; and the above proposition—or at least its proof—shows that \mathcal{T}_G is the trace of \mathcal{T}_F on G . This suggests that there may be a single σ -field \mathcal{T} on \mathcal{K} , such that \mathcal{T}_F is, for each F , the restriction of \mathcal{T} to F . (\tilde{b}) would then become a straight measurability assumption with respect to such a \mathcal{T} . A description of such a \mathcal{T} by generators would be most helpful. This needs further investigation, to fully clarify the meaning of the measurability requirement on the action sets.

3. Strategies, Payoffs and Equilibria

3.1. HISTORIES AND STRATEGIES

By a t -stage history, we mean a sequence $h_t = (g_0, g_1, \dots, g_{t-1}) \in H_t$; $\tilde{h}_t \in \tilde{H}_t$ will denote a sequence (h_t, ω_t) (with $h_t \in H_t, H_0 = \{\emptyset\}$).

$$(H_t, \mathcal{H}_t) = \left(\prod_{i=0}^{t-1} G, \bigotimes_{i=0}^{t-1} \mathcal{G} \right), \quad \text{and} \quad (\tilde{H}_t, \tilde{\mathcal{H}}_t) = (H_t \times \Omega, \mathcal{H}_t \otimes \mathcal{A}).$$

The σ -fields \mathcal{H}_t and $\tilde{\mathcal{H}}_t$ will also be viewed as sub- σ -fields of the σ -fields $\mathcal{H}_\infty = \tilde{\mathcal{H}}_\infty$ on the space $H_\infty = \tilde{H}_\infty$ of infinite histories. We will denote the disjoint union of all spaces (H_t, \mathcal{H}_t) and $(\tilde{H}_t, \tilde{\mathcal{H}}_t)$ ($t = 0, 1, 2, \dots$) by (H, \mathcal{H}) ; this is the space of all finite histories; $(\tilde{H}, \tilde{\mathcal{H}})$ will similarly denote the disjoint union of all spaces $(\tilde{H}_t, \tilde{\mathcal{H}}_t)$ ($t = 0, 1, 2, \dots$).

A (behavioral) strategy σ of player n is a transition probability from $(\tilde{H}, \tilde{\mathcal{H}})$ to $(\bar{S}_n, \mathcal{S}_n)$ assigning probability one to $S_n(\omega)$.

A pure strategy is a (behavioral) strategy where all probabilities are point masses (zero-one measures).

3.2. ASSUMPTIONS

- a) For any strategy vector σ , any player n , any initial state ω , and any stage t , the expectation $u_n^t(\sigma, \omega)$ of his payoff at stage t exists. Let $\bar{u}_n^t(\sigma, \omega) = \sum_{s < t} u_n^s(\sigma, \omega)$.
- b) For any pure strategy vector σ , any player n , and any initial state ω ,

$$v_n^\sigma(\omega) = \lim_{t \rightarrow \infty} \bar{u}_n^t(\sigma, \omega) \quad \text{exists.}$$

c)

1. Let $K(\omega) = \sup_{t, n, \sigma} |\bar{u}_n^t(\sigma, \omega)|$ where σ ranges over the pure strategies. Then $K(\omega) < \infty$ for all ω .
 2. $\int K(\omega) dq(\omega) < \infty$ for all probability distributions $q(\cdot) = p(\cdot \mid \tilde{\omega}, s)$. (In fact, it will be shown that $K(\omega)$ is measurable; in the meantime, the integral can be interpreted as a lower integral.)
- d) For every initial state, every pure strategy vector σ , and any $\epsilon > 0$, there exists t_0 such that, for any pure strategy vector τ which coincides with σ up to t_0 , $|v_n^\sigma(\omega) - v_n^\tau(\omega)|$ (continuity, using the product of the discrete topologies on the pure strategy space).
 - e) (This would follow from requirement (f) (Section 2.2.1) if the v_n^σ were uniformly bounded): for any pure strategy vector σ , and any initial state ω_0 , the mapping $s \rightarrow \int v_n^\sigma(\omega) dp(\omega \mid s, \omega_0)$ is continuous.

3.3. PAYOFFS

For technical reasons, we will also need the following concepts. A mixture of behavioral strategies (resp. a mixed strategy) of player n^1 is similarly described by an auxiliary probability space $(X_n, \mathcal{X}_n, Q_n)$, and a transition probability (resp. zero-one-valued) from $(X_n \times \bar{H}, \mathcal{X}_n \otimes \mathcal{H})$ to \bar{S}_n, \mathcal{S}_n .

Any strategy vector (of whatever type) σ , together with an initial distribution μ on Ω , induces a unique probability distribution on \mathcal{H}_∞ (such that the expectation of any positive function that depends only on the first t coordinates can be computed backwards using the strategies in the obvious way—Ionescu–Tulcea theorem). The corresponding expectation will be denoted by E_σ^μ , and by E_σ^ω when μ is a point mass at some initial state ω .

Given any t -stage history h , and any behavioral strategy σ , the conditional strategy σ^h is defined by $\sigma^h(\tilde{h}) = \sigma(h, \tilde{h})$. When σ is a strategy vector, the corresponding expectation operator $E_{\sigma^h}^\mu$ will also be denoted by

¹Our treatment of mixtures follows [2]; it has to be adapted among other reasons because the payoff function is defined directly on strategies, and does not necessarily stem from a payoff function defined on histories. We need the mixtures, even just to obtain the relevant properties of behavioral strategies.

E_σ^h when $t \geq 1$ and $\mu(\cdot) = p(\cdot \mid g_{t-1})$ (where g_{t-1} is the last element of h), and by $E_\sigma^{\tilde{h}}$ for $\tilde{h} = (h, \omega_t)$ when μ is the unit mass at ω_t (this extends the notation E_σ^ω).

Similarly, we will use the notation $\sigma^{\tilde{h}}$ where $\tilde{h} = (h, \omega)$, for the mapping $p \rightarrow \sigma^h(\omega, p)$, for all p such that $(\omega, p) \in \tilde{H}$.

Lemma 2 *Given a behavioral strategy σ :*

- a) σ^h is a strategy.
- b) σ^h is measurable in h in the sense that, for any strategy vector σ and any positive, measurable function f on $(H_\infty, \mathcal{H}_\infty)$, $E_\sigma^h f(h, \cdot)$ and $E_\sigma^{\tilde{h}} f(\tilde{h}, \cdot)$ are \mathcal{H} and $\tilde{\mathcal{H}}$ -measurable (w.r.t. $h \in H$ and $\tilde{h} \in \tilde{H}$ resp.).
- c) For any initial distribution μ and any t , those expectations are versions of $E_\sigma^\mu(f \mid \mathcal{H}_t)$ and of $E_\sigma^\mu(f \mid \tilde{\mathcal{H}}_t)$ respectively.
- d) For a strategy σ of whatever type $E_\sigma^\omega(f)$ is measurable for any positive \mathcal{H}_∞ -measurable f .

Proof. The proof is standard.

Lemma 3 *For any mixture of behavioral strategies, there exists an equivalent mixed strategy, in the sense that, whatever the strategies (in any sense) of the other players are, and for any initial distribution μ , the probability distributions induced on the space of histories are the same.*

Proof. Since \bar{S}_n is compact metric, there exist continuous mappings ϕ_n from the Cantor set $\{0, 1\}^\infty$ (with its usual embedding into $[0, 1]$) onto \bar{S}_n . For these there exist Borel-measurable selections ψ_n , e.g., $\psi_n(s) = \min\{x \mid \phi_n(x) = s\}$. Hence, given our mixture σ , we can use ψ_n to consider the transition probabilities to be transition probabilities to $[0, 1]$ (whose support is in the compact inverse image of $S_n(\omega)$): these are uniquely described by their cumulative distribution function $F : [0, 1] \rightarrow [0, 1]$ when F is right-continuous, nondecreasing and $F(1) = 1$. Now make the product of the space $(X_n, \mathcal{X}_n, \mathcal{Q}_n)$ with an infinite product of copies of $[0, 1]$ with Lebesgue measure. Given a point $(x, y_0, y_1, \dots, y_t, \dots)$ in this product, let player n at stage t play $\phi_n[\min\{z \mid F_{x, \tilde{h}_t}(z) \geq y_t\}]$: this is the associated mixed strategy, and it is easy to check that it has the required properties.

Lemma 4 *For any strategy vector and any initial state, there exists a vector of behavioral strategies that induces the same probability distribution on $(H_\infty, \mathcal{H}_\infty)$. (Those elements of the given strategy vector that happen to be behavioral strategies do not have to be changed.)*

Proof. At each stage t , the given probability distribution on histories induces a joint distribution on $(\tilde{H}_t, \tilde{\mathcal{H}}_t) \otimes (\bar{S}_n, \bar{\mathcal{S}}_n)_t$. Let ϕ_i be a countable dense subset of continuous functions on the compact metric space \bar{S}_n , containing the constant function one, and forming a vector space over the

rational. Let f_i be an $\tilde{\mathcal{H}}_t$ -measurable version of the conditional expectation of ϕ_i given $\tilde{\mathcal{H}}_t$. Let $N \in \tilde{\mathcal{H}}_t$ be the null set where either for some i, j, k , $f_i + f_j \neq f_k$ while $\phi_i + \phi_j = \phi_k$, or $f_i \neq \alpha f_j$ while $\phi_i = \alpha \phi_j$, or $f_i < 0$ while $\phi_i > 0$, or $f_i \neq 1$ while $\phi_i = 1$. For $h \notin N$, the mapping $\phi_i \rightarrow f_i(h)$ extends to a positive linear functional of norm 1 or $C(\bar{S}_n)$, hence by Riesz's theorem to a probability measure on \bar{S}_n . Thus we get a transition probability from the complement of N to $(\bar{S}_n, \bar{\mathcal{S}}_n)$, such that the conditional expectation of any continuous function on \bar{S}_n given $\tilde{\mathcal{H}}_t$ is correctly described by this transition probability. Hence this is true for any nonnegative Borel function on \bar{S}_n , hence for products of such functions with nonnegative $\tilde{\mathcal{H}}_t$ -measurable functions, hence for any nonnegative $\mathcal{H}_t \otimes S_n$ -measurable function. In particular, the conditional probability of $\{(\tilde{h}, s) \mid \tilde{h} = (h, \omega_t), s \in S_n(\omega_t)\}$ is $\tilde{\mathcal{H}}_t$ -measurable, and a.e. equal to one. Add the set where it differs from 1 to the null set N , and define the conditional probability on N by any fixed \mathcal{H} -measurable selection from the graph of S_n (Section 2, Corollary 2). This conditional probability now defines the t^{th} -stage component of player n 's behavioral strategy. Clearly this construction has the required properties.

Lemma 5 *For any initial state ω , and any vector σ of strategies (of whatever type):*

- a) *the expectations $u_n^t(\sigma, \omega) = E_\sigma^\omega$ (player n 's payoff at stage t) exist (and are finite);*
- b) *$u_n^t(\sigma, \omega)$ is measurable in ω ;*
- c) *the partial sums $\bar{u}_n^t(\sigma, \omega) = \sum_{s \leq t} u_n^s(\sigma, \omega)$ satisfy $|\bar{u}_n^t(\sigma, \omega)| \leq K(\omega)$;*
- d) *$\bar{u}_n^t(\sigma, \omega)$ converges pointwise ($t \rightarrow \infty$), say to $v_n^\sigma(\omega)$.*

Proof. By Lemma 3, it is sufficient to consider vectors of mixed strategies and, for points a) and d), we can even use Lemma 4 to consider only behavioral strategies—since payoffs at stage t are measurable functions on the space of histories.

Point a) follows then from Assumption (a) sub B, and point b) from Lemma 2

For c), we consider the product of the space of histories with the spaces $(X_n, \mathcal{X}_n, \mathcal{Q}_n)$ of each player—considering this as part of this initial state. Formally, take the initial state to be in the product of $(X, \mathcal{X}, \mathcal{Q}) = \Pi_n(X_n, \mathcal{X}_n, \mathcal{Q}_n)$ with a copy of (Ω, \mathcal{A}) . On this new state space, we have pure strategies; and if σ_n is player n 's, then $\forall x \in X_n, \sigma_n^x$ is a pure strategy of the original game (Lemma 2).

Since, by (a), player n 's payoff at stage t is integrable on this enlarged space of histories, we can compute its expectation by first taking the conditional expectation given $(\omega, x_1, x_2, \dots)$: this is (Lemma 2.c)) equal to $u_n^t(\sigma^x, \omega)$, which is (Lemma 2.b)) measurable in (ω, x) .

Hence the partial sums $\sum_{s < t} u_n^s(\sigma^x, \omega)$ are, by Assumption c), bounded in absolute value by the constant $K(\omega)$, and converge pointwise to $v_n^{\sigma^x}(\omega)$ by Assumption b). Integrating over x_1, \dots, x_n yields the partial sums $\bar{u}_n^t(\sigma, \omega)$ —which are therefore also bounded by $K(\omega)$, thus establishing c); by the dominated convergence theorem, they will converge to $\int v_n^{\sigma^x}(\omega) dQ(x)$. This establishes d).

Note that we have also shown that $v^\sigma(\omega) = \int v^{\sigma^x}(\omega) dQ(x)$.

We can now define the payoff to arbitrary strategies:

Definition 1 $v_n^\sigma(\omega)$ is the payoff to player n for the initial state ω resulting from the strategy vector σ (of whatever type).

We will also use the notations $v^\sigma(\tilde{h})$, with $\tilde{h} = (h, \omega)$, for $v^{\sigma^h}(\omega)$; and $v^\sigma(h)$, with $h = (\tilde{h}, s)$, and $\tilde{h} = (h', \omega_t)$, $s \in S(\omega_t)$, for $v^\tau(\tilde{h})$, where $\tau = \tau(s) = \sigma$ everywhere, except at stage t , where an arbitrary strategy τ_t is used, such that $[\tau_t(\tilde{h})](\{s\}) = 1$. (By Section 2.2.1.b, $\{\omega \mid s \in S(\omega)\} \in \mathcal{A}$; let $\tau_t(\omega) = s$ on this set, and be an arbitrary measurable selection (Section 2, Corollary 2) on the complement.) Clearly $v^\sigma(h)$ is well defined, and independent of the choice of τ_t .

Lemma 6 If $h \in H_t$ and $\tilde{h} \in \tilde{H}_t$, then

- a) $v^\sigma(h)$ and $v^\sigma(\tilde{h})$ depend only on the components of σ after stage t .
- b) $v^\sigma(h)$ and $v^\sigma(\tilde{h})$ are \mathcal{H} - and $\tilde{\mathcal{H}}$ -measurable.
- c) $v^\sigma(h) = \lim_{T \rightarrow \infty} \sum_{s < T} E_\sigma^h u^s$ and $v^\sigma(\tilde{h}) = \lim_{T \rightarrow \infty} \sum_{s < T} E_\sigma^{\tilde{h}} u^s$ where u^s denotes the payoff vector at stage s . The notation means that all expectations written exist, are absolutely convergent, and that the relevant limits exist and are finite.
- d)

$$1. v^\sigma(h) = u(\omega, s) + \int v^\sigma((h, \tilde{\omega})) dp(\tilde{\omega} \mid \omega, s) \text{ where } h = (h', \omega, s),$$

and

$$2. v^\sigma(\tilde{h}) = \int v^\sigma(\tilde{h}, s) d\sigma_1(s_1 \mid \tilde{h}) d\sigma_2(s_2 \mid \tilde{h}) \dots$$

All integrals are again absolutely convergent.

Proof. (a) is obvious.

(c) follows from the definitions and from Lemma 5.

(b) follows from c) and from Lemma 2(b).

(d) The first formula follows because $v^\sigma(h, \tilde{\omega}) = \lim_T \bar{u}^T(\sigma^h, \tilde{\omega})$ (by c)), and since $\|\bar{u}^T(\sigma^h, \tilde{\omega})\| \leq K(\tilde{\omega})$ (Lemma 5.c), which is $dp(\tilde{\omega} \mid \omega, s)$ -integrable for each ω, s (Assumption c), we have, by the dominated convergence theorem,

$$\int v^\sigma(h, \tilde{\omega}) dp(\tilde{\omega} \mid \omega, s) = \lim_T \int \bar{u}^T(\sigma^h, \tilde{\omega}) dp(\tilde{\omega} \mid \omega, s) = \sum_{s \geq 1} E_\sigma^h(u^s).$$

There remains to add $E_\sigma^h(u^0) = u(\omega, s)$ and to apply c) once more.

The second formula follows because, by c),

$$v^\sigma(\tilde{h}, s) = \lim_T E_\sigma^{\tilde{h}, s} \sum_{t < T} u^t = \lim_T \bar{u}^T(\tau^h(s), \omega), \text{ for } \tilde{h} = (h, \omega)$$

and $|\bar{u}^T(\tau^h(s), \omega)| \leq K(\omega)$ by Lemma 5.c), for all $s \in S(\omega)$ and all T .

Hence the dominated convergence theorem implies that

$$\int v^\sigma(\tilde{h}, s) d\sigma(s \mid \tilde{h}) = \lim_T \int \bar{u}^T(\tau^h(s), \omega) d\sigma(s \mid \tilde{h}) = \sum_t E_\sigma^{\tilde{h}}(u^t),$$

from where the result follows by another application of c).

3.4. EQUILIBRIA

Since the payoff function is now well defined for arbitrary vectors of behavioral strategies, the concept of Nash equilibrium (in behaviour strategies) is also well defined. However, to show that those qualify unambiguously as Nash equilibria of the stochastic game, we want to show that, at such equilibria, no player has profitable replies, even when not restricted in the replies by the same stringent measurability requirements as in the strategies.

Lemma 7 *a) A strategy vector σ (of whatever type) is a Nash equilibrium iff no player has a profitable pure strategy deviation.*

b) In that case, no player n has a profitable deviation even if the measurability requirement on his strategies was weakened to require only that, at each stage t , his transition probability to $S_n(\omega_t)$ be μ_t -measurable w.r.t. (x, \tilde{h}_t) , where μ_t is the probability induced on $(X_n \times \tilde{H}_t, \mathcal{X}_n \otimes \tilde{\mathcal{H}}_t)$ by the initial state, the other players' strategies, Q_n , and his own past transition probabilities.

c) $\sup_\sigma v_n^\sigma(\omega)$ and $\inf_\sigma v_n^\sigma(\omega)$ can be equivalently computed over pure strategies σ only.

Proof. a) By Lemma 3, we can take all strategies to be mixed. Assume player 1 had a profitable deviation, say τ_1 (with $(X_1\mathcal{X}_1, Q_1)$): then

$$v_1^\sigma(\omega) < \int \left[\int v_1^{\tau_1^{x_1}, \sigma_2^{x_2}, \sigma_3^{x_3}, \dots}(\omega) dQ_2(x_2) dQ_3(x_3), \dots \right] dQ_1(x_1)$$

by our formula $v^\sigma(\omega) = \int v^{\sigma^x}(\omega) dQ(x)$. But then there exists some x_1 with

$$v_1^\sigma(\omega) < \int v_1^{\tau_1^{x_1}, \sigma_2^{x_2}, \sigma_3^{x_3}, \dots}(\omega) dQ_2(x_2) dQ_3(x_3), \dots = v_1^{\tau_1^{x_1}, \sigma_2, \sigma_3, \dots}(\omega)$$

(by the same formula), and $\tau_1^{x_1}$ is a pure strategy.

b) Since \tilde{S}_n is compact metric, the σ -field \mathcal{S}_n is separable—hence is generated by a countable, dense Boolean subalgebra. For any $S \in \mathcal{S}_n$, its probability is μ_t -measurable—hence there exists an $\mathcal{X}_n \otimes \tilde{\mathcal{H}}_t$ -measurable set of μ_t -probability zero, such that on the complement of this set, the probability of S is $\mathcal{X}_n \otimes \tilde{\mathcal{H}}_t$ -measurable. Take the union of those null sets when S varies over the countable subalgebra, and redefine there the transition probability to be a point mass at some measurable selection from $S_n(\omega)$ (Section 2, Corollary 2). We now have a true strategy of player 1; since we changed only on null sets, it is obvious by induction on t the μ_t 's did not change—hence in particular the induced probability distribution on histories—and thus the payoff—will be the same.

c) One changes each player's strategy in turn to a pure strategy, each time using the same argument as in the proof of a).

Remarks. 1) As seen from the proof of Lemma 4, we could have similarly weakened the countable additivity requirement on player n 's strategy in (b), and only required that for each $S \in \mathcal{S}_n$, some μ_t -measurable function P_S be given, s.t. (α) $P_S \geq 0$, μ_t a.e., (β) $P_{\tilde{S}_n} = 1$ a.e., (γ) for any disjoint sequence $S_i \in \mathcal{S}_n$, $P_{\cup S_i} = \sum_i P_{S_i}$ μ_t -a.e. and (δ) $\forall A \in \mathcal{X}_n \otimes \tilde{\mathcal{H}}_t, \forall S \in \mathcal{S}_n$, if $S \cap S_n(\omega_t) = \emptyset$ for all ω_t in the projection of A , then $P_S = 0$ μ_t -a.e. on A . (I.e., we can just ask for a probability on \tilde{S}_n with values in $L_\infty(\mu_t)$ with the weak*-topology.)

2) Remark 1 implies that player 1 has no profitable reply in any sense, i.e., in any sense of reply for which the induced probability distribution on histories satisfies that (a) at each stage t , the pure strategy choices s_n of the different players are conditionally independent given $\tilde{\mathcal{H}}_t$; (b) those conditional distributions for the other players are described by their strategies; (c) $\text{Prob}(s_{n,t} \in S_n(\omega_t)) = 1$ for each t and n ; (d) ω_0 is the given initial state with probability one; and (e) for all t , the conditional distribution of ω_t given h_t is described by the transition probabilities of the stochastic

game. Indeed, any such probability distribution can be described as arising from some reply of player 1 in the sense of Remark 1.

For example, player 1 could use a sequence of spaces (X^i, \mathcal{X}^i) , and at each stage t , use a transition probability from the product of \bar{H}_t and the previous X^i 's to $X^t \times \bar{S}_1 \dots$ (And the conditions on the transition probabilities could be weakened, as in Remark 1.) Such “strategies” do not fit under our previous definitions, yet conditions (a)-(e) are satisfied.

3) Similarly, we can now see that the game is well defined, whatever strategy spaces are considered: for whatever concept of strategy, a strategy vector and an initial state will induce a probability distribution on the space of histories satisfying (a), (c), (d) and (e). One could then use the previous construction for all players in turn, to construct a behavioral strategy vector that induces the same probability distribution on the space of histories. Therefore, for any strategy n^{tuple} of whatever type, the expected payoffs \bar{u}_n^t for all t and n will exist and be finite, and will form a summable series, thus defining the payoff. In particular, with arbitrary strategy spaces (containing the behavioral strategies) the set of feasible payoff vectors is the same as with just behavioral strategies, and Nash equilibria of the game restricted to behavioral strategies are still Nash equilibria on the arbitrary strategy spaces.

4) Lemma 7 and the above remarks imply that the Nash equilibria of the game restricted to behavioral strategies are completely unambiguously “Nash equilibria of the stochastic game.” Thus we consider henceforth only behavioral strategies.

3.5. A REFORMULATION: BASIC PROPERTIES OF THE MODEL

Assumptions (a), (b) and (c) were used basically just to show that the payoff function $v^\sigma(\omega)$ was unambiguously defined and had the intended meaning, and satisfied the recursion formula in Lemma 6.d—and that the concept of Nash equilibrium of the game restricted to behavioral strategies was perfectly satisfactory.

One might be able to reach the same conclusions (maybe with another intended meaning for $v^\sigma(\omega)$) in many models where those assumptions are not satisfied. We now reformulate our model so that the results of this paper will then still apply. This may be seen as a first step towards extending the theorem to cover payoff functions which are arbitrary functions of histories, instead of just additively separable ones (similar recursion equations are obtained in those cases, when including the “past payoff” in the $v^\sigma(h)$).

We take as primitive datum upper bounds $\bar{v}_n^\sigma(\omega)$ and lower bounds $\underline{v}_n^\sigma(\omega)$ (possibly infinite) for the payoff to player n resulting from the strategy vector σ and the initial state ω .

We assume that they depend only on $\sigma(\omega)$, and that $\underline{v}_n^\sigma(\omega) \leq \bar{v}_n^\sigma(\omega)$. We write $v_n^\sigma(\omega)$ whenever $\underline{v}_n^\sigma(\omega) = \bar{v}_n^\sigma(\omega)$. We say that v_n^σ exists if v_n^σ exists for all ω and is measurable. We say that v^σ exists if v_n^σ exists for all n .

Proposition 3 *All statements from here on are still valid under the following assumptions (in addition to the requirements (A) of Section 2):*

- (a) *There exists a pure strategy vector σ_0 such that v^{σ_0} exists.*
- (b) *$\forall \sigma, \forall n, \forall \omega, \bar{v}_n^\sigma(\omega) \leq \sup_\tau \bar{v}_n^\tau(\omega)$ and $\underline{v}_n^\sigma(\omega) \geq \inf_\tau \underline{v}_n^\tau(\omega)$ where τ ranges over all pure strategy vectors.*
- (c) *Define $\bar{v}^\sigma(\tilde{h}) = \bar{v}^{\sigma^h}(\omega)$ for $\tilde{h} = (h, \omega)$. Then, for all $h \in H$, denoting by $\bar{\int}$ an upper integral,*

$$\bar{v}^\sigma(h, \omega) \leq \bar{\int} [u(\omega, s) + \bar{v}^\sigma(h, \omega, s, \tilde{\omega})] dp(\tilde{\omega} \mid \omega, s) d\sigma(s \mid h, \omega)$$

and similarly for $\underline{v}^\sigma(\tilde{h})$, using a lower integral.

- (d) *For any pure strategy vector τ , if σ^i is a sequence of pure strategy vectors such that v^{σ^i} exists and has finite values and such that σ^i coincides with τ during the first i stages, then v^τ exists and $v^\tau(\omega) = \lim_{i \rightarrow \infty} v^{\sigma^i}(\omega)$.*
- (e) *For any pure strategy vector σ for which v^σ exists, and any initial state ω , $\int v_n^\sigma(\tilde{\omega}) dp(\tilde{\omega} \mid \omega, s)$ exists for all $s \in S(\omega)$ and is real-valued and continuous in s for all n .*

Note that (a) is not an assumption of existence of pure strategies; we observed earlier that this already follows from Section 2.

Under our previous assumptions, we could define $\bar{v} = \underline{v} = v$ for all σ and ω ; the validity of (a) stems from Lemma 6.b, of (b) from Lemma 7.c and of (c) from Lemma 6.d. (d) and (e) follow immediately from the corresponding assumptions in Section 3.3.2.

Note finally that we can still use Lemma 2 in the sequel, as it depends only on the data of the model.

Lemma 8 (a) *v^σ exist for any pure strategy vector σ ;*

- (b) *there exist (stationary) pure strategy vectors $\bar{\sigma}_n$ and $\underline{\sigma}_n$ such that, letting $\bar{w}_n(\omega) = v_n^{\bar{\sigma}_n}(\omega)$ and $\underline{w}_n(\omega) = v_n^{\underline{\sigma}_n}(\omega)$, \bar{w}_n and \underline{w}_n are real-valued and for any pure strategy vector σ , $\underline{w}_n \leq v_n^\sigma \leq \bar{w}_n$.*

Proof. We first show that v^τ exists and is real-valued for any stationary pure strategy $\tau = (\tau_0, \tau_0, \tau_0, \dots)$.

Let $\sigma_{k+1} = (\tau_0, \sigma_k)$, with σ_0 as in Proposition 3.a. We show by induction that v^{σ_k} exists: for $k = 0$, this follows from (3.a). By (3.e), $\int v^{\sigma_k}(\tilde{\omega}) dp(\tilde{\omega} \mid \omega, s)$ exists (and is finite) for all (ω, s) —in particular for $s = \tau_0(\omega)$. Since

p is a transition probability and since v^{σ_k} is measurable, the existence of all integrals implies that the integral is measurable. Therefore, by (3.c), $\bar{v}_n^{\sigma_{k+1}}(\omega) \leq u_n(\omega, \tau_0(\omega)) + \int v_n^{\sigma_k}(\tilde{\omega}) dp(\tilde{\omega} \mid \omega, \tau_0(\omega)) \leq \underline{v}_n^{\sigma_{k+1}}(\omega)$. Since $\underline{v} \leq \bar{v}$, we have for all ω :

$$v^{\sigma_{k+1}}(\omega) = u(\omega, \tau_0(\omega)) + \int v^{\sigma_k}(\tilde{\omega}) dp(\tilde{\omega} \mid \omega, \tau_0(\omega)).$$

Measurability of $v^{\sigma_{k+1}}$ follows, by composition: $v^{\sigma_{k+1}}$ exists and is finite-valued. (7.d) implies then the existence of v^τ . Since $\tau = (\tau_0, \tau)$, repeating once more the above argument then implies that v^τ is finite-valued.

Note that we have shown the following:

Claim 1 *If σ is a pure strategy vector such that v^σ exists, and τ_0 a pure strategy vector in the one-shot game, then $v^{(\tau_0, \sigma)}$ exists, has finite values and satisfies $v^{(\tau_0, \sigma)}(\omega) = u(\omega, \tau_0(\omega)) + \int v^\sigma(\tilde{\omega}) dp(\tilde{\omega} \mid \omega, \tau_0(\omega))$.*

We now start proving (b), assuming $n = 1$ and proving the existence of $\bar{\sigma}_1$ and \bar{w}_1 . The proof for the other cases is identical.

Fix σ_0 to be some stationary pure strategy vector. Assume that σ_m is some pure strategy vector for which v^{σ_m} exists and is finite-valued. Let $w_m(\omega) = v_1^{\sigma_m}(\omega)$, and define

$$w_{m+1}(\omega) = \sup_{s \in S(\omega)} [u_1(\omega, s) + \int w_m(\tilde{\omega}) dp(\tilde{\omega} \mid \omega, s)].$$

By (e), the integral exists for all s —so, as remarked earlier, the bracketed term is measurable in (ω, s) , and real-valued and continuous in s for each ω . Thus the supremum is achieved, and is finite. The continuity implies also that measurability of w_{m+1} follows from Section 2, Corollary 2.

Thus w_{m+1} is real-valued, measurable, and satisfies

$$w_{m+1}(\omega) = \max_{s \in S(\omega)} [u_1(\omega, s) + \int w_m(\tilde{\omega}) dp(\tilde{\omega} \mid \omega, s)].$$

Let $X(\omega) = \{(s, x) \mid s \in S(\omega), x = u_1(\omega, s) + \int w_m(\tilde{\omega}) dp(\tilde{\omega} \mid \omega, s)\}$. By Section 2, Corollary 1, X is a measurable map to $\mathcal{K}_{\mathbb{R} \times \bar{S}}^*$. Similarly, $Y(\omega) = \{(s, x) \mid s \in S(\omega), x = w_{m+1}(\omega)\}$ is measurable, thus so is $T(\omega) = \{s \in S(\omega) \mid w_{m+1}(\omega) = u_1(\omega, s) + \int w_m(\tilde{\omega}) dp(\tilde{\omega} \mid \omega, s)\}$ —the projection on \bar{S} of $X(\omega) \cap Y(\omega)$, by [5], Proposition 6, (d) and (e). Hence by [5], Proposition 7(c), T has a measurable selection; i.e., there exists a strategy vector $\tau(\omega)$ of the one-shot game, such that $w_{m+1}(\omega) = u_1(\omega, \tau(\omega)) + \int w_m(\tilde{\omega}) dp(\tilde{\omega} \mid \omega, \tau(\omega))$. Let $\sigma_{m+1} = (\tau, \sigma_m) : \sigma_{m+1}$ is a pure strategy vector, and, by Claim 1, $v^{\sigma_{m+1}}$ exists and $w_{m+1}(\omega) = v_1^{\sigma_{m+1}}(\omega)$.

This completes the inductive definition of the σ_m 's and the w_m 's. Note that $w_1 \geq w_0$, since the choice of s repeating the same stationary σ_0 was available in the maximization, and would have yielded w_0 . Therefore, by induction on m , $w_{m+1} \geq w_m$.

Now let $w = \lim_{m \rightarrow \infty} w_m$: w is measurable, $w(\omega) > -\infty$ everywhere, and

$$\forall \omega, s, m \quad w(\omega) \geq w_{m+1}(\omega) \geq u_1(\omega, s) + \int w_m(\tilde{\omega}) d p(\tilde{\omega} \mid \omega, s).$$

Since w_m is integrable and increasing with m , the monotone convergence theorem implies that $w(\omega) \geq u_1(\omega, s) + \int w(\tilde{\omega}) d p(\tilde{\omega} \mid \omega, s)$; hence

$$w(\omega) \geq \sup_{s \in S(\omega)} \left[u_1(\omega, s) + \int w(\tilde{\omega}) d p(\tilde{\omega} \mid \omega, s) \right]$$

(and all integrals in the right-hand member are well defined).

Further, the right-hand member is $\geq w_{m+1}(\omega)$, since $w \geq w_m$. Thus

$$w(\omega) = \sup_{s \in S(\omega)} \left[u_1(\omega, s) + \int w(\tilde{\omega}) d p(\tilde{\omega} \mid \omega, s) \right] \quad \text{for all } \omega.$$

We first show that the supremum is attained. Denote the bracketed term by $f(\omega, s)$. Fix $\epsilon > 0$ and $K > 0$, and let $\varphi(\omega) = K$ if $w(\omega) = +\infty$, $\varphi(\omega) = w(\omega) - \epsilon$ elsewhere: φ is a real-valued measurable function. Let $N(\omega) = \min\{m \mid w_m(\omega) \geq \varphi(\omega)\}$: N is measurable and integer-valued.

Define the pure strategy vector σ by $\sigma(\omega) = \sigma_{N(\omega)}(\omega)$. By our assumption that $\bar{v}_n^\sigma(\omega)$ and $v_n^\sigma(\omega)$ depend only on $\sigma(\omega)$, we have for all ω and all n , $v_n^\sigma(\omega) = v_n^{\sigma_{N(\omega)}}(\omega)$, which is measurable and real-valued: v^σ exists and is real-valued, and $w(\omega) \geq v_1^\sigma(\omega) = w_{N(\omega)}(\omega) \geq \varphi(\omega)$.

Thus, again using (7.e), we have that

$$(\omega, s) \rightarrow u_1(\omega, s) + \int v_1^\sigma(\tilde{\omega}) d p(\tilde{\omega} \mid \omega, s)$$

is measurable, is continuous and real-valued in s for each ω , and that the integrals are absolutely convergent.

Going to the (uniform) limit in ϵ ($\epsilon \rightarrow 0$), we obtain that

$$\psi : (\omega, s) \rightarrow u_1(\omega, s) + \int_{w < +\infty} w(\tilde{\omega}) d p(\tilde{\omega} \mid \omega, s) + \int_{w = +\infty} v_1^\sigma(\tilde{\omega}) d p(\tilde{\omega} \mid \omega, s)$$

has the same properties.

And $f(\omega, s) = \psi(\omega, s) + I(\omega, s)$, where $I(\omega, s)$ is $(+\infty)$ times the indicator function of

$$\tilde{A} = \{(\omega, s) \mid p(A \mid \omega, s) > 0\} \quad \text{with} \quad A = \{\omega \mid w(\omega) = +\infty\}.$$

$A \in \mathcal{A}$, hence $\tilde{A} \in \mathcal{G}$, and the norm continuity in s of $p(\cdot \mid \omega, s)$ implies that $\{s \mid (\omega, s) \in \tilde{A}\}$ is open.

Thus, for each ω , the maximum is achieved—for $\omega \in A$, at any s such that $(\omega, s) \in \tilde{A}$, and for $\omega \notin A$, at any s maximizing $\psi(\omega, s)$. Let $h(\omega) = \max_{s \in S(\omega)} p(A \mid \omega, s)$: h is real-valued, and measurable by Section 2, Corollary 2, as before. Let again, for $\omega \in A$, $X(\omega) = \{(s, x) \mid s \in S(\omega), x = p(A \mid \omega, s)\}$, $Y(\omega) = \{(s, x) \mid s \in S(\omega), x = h(\omega)\}$ and for $\omega \notin A$, $X(\omega) = \{(s, x) \mid s \in S(\omega), x = \psi(\omega, s)\}$, $Y(\omega) = \{(s, x) \mid s \in S(\omega), x = w(\omega)\}$. As before, X and Y are measurable maps to $\mathcal{K}_{\bar{S} \times \mathbb{R}}^*$, hence so is the projection of \bar{S} of their intersection. This admits therefore a measurable selection, which is a pure strategy vector of the one-shot game, τ_0 , at which $f(\omega, s)$ achieves its maximum: for all ω

$$w(\omega) = u_1(\omega, \tau_0(\omega)) + \int w(\tilde{\omega}) \, d p(\tilde{\omega} \mid \omega, \tau_0(\omega)).$$

Now consider a sequence of pure strategy vectors σ_m , like those we defined above, such that v^{σ_m} exists and is finite-valued for all m , and such that $v_1^{\sigma_m}$ increases to w . By our above claim, the $v^{(\tau_0, \sigma_m)}$ will exist and have finite values, and will satisfy

$$v_1^{(\tau_0, \sigma_m)}(\omega) = u_1(\omega, \tau_0(\omega)) + \int v_1^{\sigma_m}(\tilde{\omega}) \, d p(\tilde{\omega} \mid \omega, \tau_0(\omega)).$$

Thus, by the monotone convergence theorem, $v_1^{(\tau_0, \sigma_m)}(\omega)$ will increase to

$$w(\omega) = u_1(\omega, \tau_0(\omega)) + \int w(\tilde{\omega}) \, d p(\tilde{\omega} \mid \omega, \tau_0(\omega)).$$

Thus the sequence (τ_0, σ_m) has the same properties as the original sequence σ_m . Now let $\sigma_{m,0} = \sigma_m, \sigma_{m,k+1} = (\tau_0, \sigma_{m,k})$: for each k , $v^{\sigma_{m,k}}$ exists and has finite values, and when $m \rightarrow \infty$, $v_1^{\sigma_{m,k}}$ increases to $w(\omega)$.

Let $\bar{\tau}$ denote the stationary pure strategy (τ_0, τ_0, \dots) . Now fix an initial state ω , and choose m_k such that, writing for short τ_k for $\sigma_{m_k,k}$, we have $v_1^{\tau_k}(\omega) \geq w(\omega) - k^{-1}$ if $w(\omega) < +\infty$, and $v_1^{\tau_k}(\omega) \geq k$ if $w(\omega) = +\infty$: the τ_k are a sequence of pure strategy vectors for which v^{τ_k} exists and is finite, and which coincide for the first k stages with the pure strategy vector $\bar{\tau}$: by (7.d), we have $w(\omega) = \lim_{k \rightarrow \infty} v_1^{\tau_k}(\omega) = v_1^{\bar{\tau}}(\omega)$.

Since $\bar{\tau}$ is a stationary pure strategy vector, $v^{\bar{\tau}}$ exists and has finite values: we have shown that there exists a stationary pure strategy vector $\bar{\tau}$, such that $v^{\bar{\tau}}$ exists and has finite values, and such that

$$v_1^{\bar{\tau}}(\omega) = \max_{s \in S(\omega)} \left[u_1(\omega, s) + \int v_1^{\bar{\tau}}(\tilde{\omega}) d p(\tilde{\omega} | \omega, s) \right]. \quad (1)$$

To finish the proof of the lemma, there remains to show that v^σ exists and is finite-valued for any pure strategy vector σ , and that it satisfies $v_1^\sigma \leq v_1^{\bar{\tau}}$. By (7.d), it is sufficient to show this for pure strategies σ that coincide with $\bar{\tau}$ after k stages. We prove this by induction on k —it is true for $k = 0$, and the induction step follows from (7.c) and the dual inequality (applying not only our equation (*) for $v_1^{\bar{\tau}}$, but the analog lower bound with a $\underline{\tau}$ and a $\min_{s \in S(\omega)}$, and the similar equations for the other players).

The next proposition should “normally” come only in Section 5; in particular, the operator Φ is central to the proof. But there is one small part of the conclusions which will be needed in Proposition 5.

We use vectors p for payoff vectors.

Proposition 4 *Let $W(\omega) = \{p \mid \forall n, \underline{w}_n(\omega) \leq p_n \leq \bar{w}_n(\omega)\}$ and denote by \mathcal{F} the set of \mathcal{A} -measurable selections from W . $\forall f \in \mathcal{F}$, let $[\Phi(f)](\omega)(s) = u(\omega, s) + \int f(\tilde{\omega}) d p(\tilde{\omega} | \omega, s) : [\Phi(f)](\omega)$ is a function on $S(\omega)$. Finally, let $K(\omega) = \{[\Phi(f)](\omega) \mid f \in \mathcal{F}\}$. Then*

- a) $K(\omega)$ is compact in the uniform topology on the space of continuous functions on $S(\omega)$;
- b) any limit point φ of a sequence $[\Phi(f_i)](\omega)$, $f_i \in \mathcal{F}$, is also the limit of a sequence $[\Phi(g_k)](\omega)$, where the g_k are measurable selections from $(\text{Lim} f_i)(\tilde{\omega}) = \{p \mid p \text{ is a limit point of } f_i(\tilde{\omega})\}$;
- c) for any finite subset of $S(\omega)$, there exists such a g such that $[\Phi(g)](\omega)$ coincides with φ on this subset.

Proof. Since $u(\omega, s)$ is just a fixed additive term, continuous in s , we can neglect it, and assume that $u = 0$. Let $w_n = \bar{w}_n - \underline{w}_n$. Let $A_k = \{\omega \mid \exists n : w_n(\omega) \geq k\} \in \mathcal{A}$. $A_k \searrow \emptyset$. Let $B_k = \Omega \setminus A_k$, and $g_k = \underline{w}_n + w_n \mathbf{I}_{A_k}$. g_k is the payoff function to the pure strategy consisting of using $\bar{\sigma}^n$ on A_k , and $\underline{\sigma}^n$ elsewhere. Thus, by (7.e), the $[\Phi(g_k)](\omega)$ are a sequence of continuous, real-valued functions on $S(\omega)$, decreasing to the continuous, real-valued function $[\Phi(\underline{w}_n)](\omega)$.

By Dini's theorem, they converge uniformly (compactness of $S(\omega)$). Hence

$$\forall \epsilon, \forall n, \exists k : \int_{A_k} w_n(\tilde{\omega}) d p(\tilde{\omega} | \omega, s) < \epsilon \quad \forall s \in S(\omega).$$

Hence

$$\forall \epsilon, \exists k : \int_{A_k} \max_n w_n(\tilde{\omega}) \, d p(\tilde{\omega} \mid \omega, s) < \epsilon \quad \forall s \in S(\omega).$$

Assume k_0 fixed, with the above property. Since the set of measures $p(d\tilde{\omega} \mid \omega, s)$ for $s \in S(\omega)$ is norm-compact (A.f), there exists a probability measure μ on (Ω, \mathcal{A}) such that the $p(d\tilde{\omega} \mid \omega, s)$ are dominated by μ , hence norm-compact in $L_1(\mu)$. Choose μ such that the w_n 's are μ -integrable. Now consider a sequence of measurable functions f_n^i satisfying $0 \leq f_n^i \leq w_n$. Since f_n^i is integrably bounded, we can extract a subsequence such as to have it weakly convergent, say to g_n . For each k and n , the sequence $f_n^i \cdot \mathbf{I}_{B_k}$ is bounded in $L_\infty(\mu)$; hence, for all k and n , $f_n^i \cdot \mathbf{I}_{B_k}$ converges $\sigma(L_\infty(\mu), L_1(\mu))$ to $g_n \cdot \mathbf{I}_{B_k}$. One can choose for g an \mathcal{A} -measurable version in its equivalence class, and since $0 \leq g_n \leq w_n$ μ -a.e., and w_n is \mathcal{A} -measurable, we can in addition select $0 \leq g_n \leq w_n$ everywhere. Now $0 \leq f_n^i \leq w_n$ implies that

$$\left| \int f_n^i(\tilde{\omega}) \, p(d\tilde{\omega} \mid \omega, s) - \int_{B_{k_0}} f_n^i(\tilde{\omega}) \, p(d\tilde{\omega} \mid \omega, s) \right| \leq \int_{A_{k_0}} w_n(\tilde{\omega}) \, p(d\tilde{\omega} \mid \omega, s) \leq \epsilon$$

and similarly for g one gets

$$\left| \int g_n(\tilde{\omega}) \, p(d\tilde{\omega} \mid \omega, s) - \int_{B_{k_0}} g_n(\tilde{\omega}) \, p(d\tilde{\omega} \mid \omega, s) \right| \leq \epsilon,$$

for all $s \in S(\omega)$ in both cases. Since $\mathbf{I}_{B_k} f_n^i$ converges weak* to $\mathbf{I}_{B_k} \cdot g_n$, it converges uniformly on compact sets of $L_1(\mu)$ (the norm topology on L_1 being the topology of uniform convergence on bounded subsets of L_∞). Thus, $\forall k, n, \forall \epsilon, \exists i_{n,k}, \forall i \geq i_{n,k}, \forall s \in S(\omega)$,

$$\left| \int_{B_k} f_n^i(\tilde{\omega}) \, p(d\tilde{\omega} \mid \omega, s) - \int_{B_k} g_n(\tilde{\omega}) \, p(d\tilde{\omega} \mid \omega, s) \right| \leq \epsilon.$$

Let $i_0 = \max_n i_{n,k_0}$: then $\forall i \geq i_0, \forall s \in S(\omega), \forall n$

$$\left| \int f_n^i(\tilde{\omega}) \, p(d\tilde{\omega} \mid \omega, s) - \int g_n(\tilde{\omega}) \, p(d\tilde{\omega} \mid \omega, s) \right| \leq 3\epsilon.$$

Thus, from any \mathcal{A} -measurable sequence f_n^i satisfying $0 \leq f_n^i \leq w_n$ we can extract a subsequence i_j converging weakly to an \mathcal{A} -measurable function f satisfying also $0 \leq f_n \leq w_n$. And then $\int f^{i_j}(\tilde{\omega}) \, p(d\tilde{\omega} \mid \omega, s)$ converges to $\int f(\tilde{\omega}) \, p(d\tilde{\omega} \mid \omega, s)$ uniformly on $S(\omega)$. Now the norm continuity of $p(d\tilde{\omega} \mid \omega, s)$ in $s \in S(\omega)$ implies that, for bounded measurable f , $\int f(\tilde{\omega}) \, p(d\tilde{\omega} \mid \omega, s)$ is continuous on $S(\omega)$. For f \mathcal{A} -measurable, $0 \leq f_n \leq w_n$, define f^i by

$f_n^i = \min(i, f_n)$. Then f^i converges weakly to f ; hence $\int f^i p(d\tilde{\omega} \mid \omega, s)$ converges uniformly to $\int f p(d\tilde{\omega} \mid \omega, s)$ on $S(\omega)$. Since the f^i 's are bounded, their integrals are continuous on $S(\omega)$, hence so is $\int f p(d\tilde{\omega} \mid \omega, s)$.

Thus the above-mentioned uniform convergence is in the space of continuous functions on $S(\omega)$.

[5], Lemma 1 implies that we can select f such that furthermore $f(\omega) \in Co(\text{Lim } f^i)(\omega)$ everywhere, and $f(\omega) \in (\text{Lim } f^i)(\omega)$ μ -a.e. on the atoms of μ .

Add now \underline{w} (the vector \underline{w}_n) to all functions f^i and f : since \underline{w} is \mathcal{A} -measurable, we get f^i and f in \mathcal{F} , and since $\int \underline{w}(\tilde{\omega}) p(d\tilde{\omega} \mid \omega, s)$ is real-valued and continuous in s (Proposition 3.e), we indeed obtain (1) that $\forall f \in \mathcal{F}$, $[\Phi(f)](\omega)$ is continuous on $S(\omega)$, and (2), that given a sequence $f^i \in \mathcal{F}$ one can find $f \in \mathcal{F}$ and a subsequence along which $[\Phi(f^i)](\omega)$ converges uniformly in the space of continuous functions on $S(\omega)$ to $[\Phi(f)](\omega)$ —hence (a) is established—and such that $f(\omega) \in \text{Convex hull}[(\text{Lim } f^i)(\omega)]$, with $f(\omega) \in [(\text{Lim } f^i)(\omega)]$, μ a.e. on the atoms of μ . By [5], Proposition 10.a, $(\text{Lim } f^i)$ is a measurable map to $\mathcal{K}^*(\mathbb{R}^\ell)$; hence by [5], Proposition 11.b, for any bounded \mathbb{R}^k -valued measure ν which is absolutely continuous w.r.t. μ , we can find a measurable selection g from $(\text{Lim } f^i)$ such that $\int g d\nu = \int f d\nu$. I.e., for any $s_1, \dots, s_k \in S(\omega)$, $[\Phi(f)](\omega)$ coincides with $[\Phi(g)](\omega)$, for some measurable selection g from $\text{Lim}(f^i)$, at all points s_1, \dots, s_k . This establishes (c). In particular, $[\Phi(f)](\omega)$ is the pointwise closure of those $[\Phi(g)](\omega)$, hence by (a) in their uniform closure. This establishes (b).

Proposition 5 (a)

1. v^σ exists for any strategy vector σ .
2. $v^\sigma(\tilde{h}_t)$ and $v^\sigma(h_t)$ are measurable.
3. $v^\sigma \in \mathcal{F}$.

(b)

1. $v^\sigma(h) = u(\omega, s) + \int v^\sigma(h, \tilde{\omega}) d p(\tilde{\omega} \mid \omega, s)$, for $h = (h', \omega, s)$, exists for any strategy vector σ , and is finite-valued and measurable.
2. Similarly, $v^\sigma(\tilde{h}) = \int v^\sigma(\tilde{h}, s) d \sigma_1(s_1 \mid \tilde{h}) d \sigma_2(s_2 \mid \tilde{h}) \dots$
3. All those integrals are absolutely convergent, even when $v^\sigma(\tilde{h})$ is expressed by one single double integral in terms of $v^\sigma(\tilde{h}, s, \tilde{\omega})$.

- (c) For any $\epsilon > 0$, there exists an integer-valued measurable function $N(\omega)$, such that for each ω , and any two strategy vectors σ and τ that coincide during the first $N(\omega)$ stages, $\|v^\sigma(\omega) - v^\tau(\omega)\| < \epsilon$.

Proof. Let $n = 1$ in Lemma 8.b, and drop the indices n : we have $\bar{w}(\omega) = v_1^{\bar{\sigma}}(\omega)$ and $\underline{w}(\omega) = v_1^{\underline{\sigma}}(\omega)$, $\bar{\sigma} = (\bar{\sigma}_0, \bar{\sigma}_0, \dots)$ and $\underline{\sigma} = (\underline{\sigma}_0, \underline{\sigma}_0, \dots)$ being two stationary pure strategies. \bar{w} and \underline{w} are real-valued and measurable, and

satisfy $\bar{w}(\omega) \geq u_1(\omega, s) + \int \bar{w}(\tilde{\omega}) d p(\tilde{\omega} \mid \omega, s)$ and $\underline{w}(\omega) \leq u_1(\omega, s) + \int \underline{w}(\tilde{\omega}) d p(\tilde{\omega} \mid \omega, s) \quad \forall s \in S(\omega)$, with equality resp. for $s = \bar{\sigma}_0(\omega)$ and for $s = \underline{\sigma}_0(\omega)$.

Subtracting both inequalities yields, with $w_0 = \bar{w} - \underline{w}$, that

$$w_0(\omega) \geq \int w_0(\tilde{\omega}) d p(\tilde{\omega} \mid \omega, s) \quad \forall s \in S(\omega).$$

w_0 is a positive, real-valued, measurable function; and Proposition 4.a implies that for any measurable function f , with $0 \leq f \leq w_0$, $\int f(\tilde{\omega}) d p(\tilde{\omega} \mid \omega, s)$ is continuous on $S(\omega)$ (since $f \in \mathcal{F}_1 - \mathcal{F}_1$). Assuming w_i defined and measurable, $0 \leq w_i \leq w_0$, let $w_{i+1} = \max_{s \in S(\omega)} \int w_i(\tilde{\omega}) d p(\tilde{\omega} \mid \omega, s)$; the maximum exists because we have seen that the integral is continuous in s . The measurability of w_{i+1} follows for the same reason—using as before Corollary 2 of Section 2.

Further, by our inequality for w_0 , one will have $w_1 \leq w_0$, and hence by induction $w_{i+1} \leq w_i$. This completes the definition of the decreasing sequence w_i .

Let $w_\infty = \lim_{i \rightarrow \infty} w_i$; w_∞ is measurable, $0 \leq w_\infty \leq w_0$, and $w_{i+1}(\omega) \geq \int w_i(\tilde{\omega}) d p(\tilde{\omega} \mid \omega, s) \geq \int w_\infty(\tilde{\omega}) d p(\tilde{\omega} \mid \omega, s)$ for all i and s implies $w_\infty(\omega) \geq \max_{s \in S(\omega)} \int w_\infty(\tilde{\omega}) d p(\tilde{\omega} \mid \omega, s)$. On the other hand, $w_\infty \leq w_{i+1}$ implies that, for all i and ω , $\{s \mid \int w_i(\tilde{\omega}) d p(\tilde{\omega} \mid \omega, s) \geq w_\infty(\omega)\}$ is nonempty, and compact by continuity. Take any s_0 in the intersection of this decreasing sequence of nonempty compact sets: we have $\int w_i(\tilde{\omega}) d p(\tilde{\omega} \mid \omega, s_0) \geq w_\infty(\tilde{\omega})$ for all i ; since the w_i 's decrease to w_∞ and the integrals are finite, the monotone convergence theorem implies that $\int w_\infty(\tilde{\omega}) d p(\tilde{\omega} \mid \omega, s_0) \geq w_\infty(\tilde{\omega})$: thus

$$w_\infty(\omega) = \max_{s \in S(\omega)} \int w_\infty(\tilde{\omega}) d p(\tilde{\omega} \mid \omega, s).$$

As in the proof of Lemma 8, this equation implies that there exists a pure strategy vector of the one-shot game τ_∞ such that

$$w_\infty(\omega) = \int w_\infty(\tilde{\omega}) d p(\tilde{\omega} \mid \omega, \tau_\infty(\omega)).$$

Denote by $\bar{\tau}_k$ (resp. $\underline{\tau}_k$) the pure strategy vector consisting in playing k times τ_∞ , then $\bar{\sigma}$ (resp. $\underline{\sigma}$). τ will denote the stationary pure strategy vector $(\tau_\infty, \tau_\infty, \dots)$. Let also $\bar{v}_k = v_1^{\bar{\tau}_k}$, $\underline{v}_k = v_1^{\underline{\tau}_k}$, $v = v_1^\tau$; those exist by Lemma 8. We have $\bar{v}_0 = \bar{w}$, $\underline{v}_0 = \underline{w}$, hence $\bar{v}_0 - \underline{v}_0 = w_0 \geq w_\infty$. Assume $\bar{v}_k - \underline{v}_k \geq w_\infty$. By (7.c),

$$\begin{aligned}
(\bar{v}_{k+1} - \underline{v}_{k+1})(\omega) &= \int [(\bar{v}_k - \underline{v}_k)](\tilde{\omega}) \, d p(\tilde{\omega} \mid \omega, \tau_\infty(\omega)) \\
&\geq \int w_\infty(\tilde{\omega}) \, d p(\tilde{\omega} \mid \omega, \tau_\infty(\omega)) = w_\infty(\omega).
\end{aligned}$$

Hence, for all k , $\bar{v}_k - \underline{v}_k \geq w_\infty$. By (7.d), for fixed ω , both \bar{v}_k and \underline{v}_k are, for k sufficiently large, ϵ -close to v ; hence $0 \leq w_\infty \leq 2\epsilon$ whatever ϵ is. Thus w_i decreases to zero.

Let $N(\omega) = \min\{i \mid w_i(\omega) < \epsilon\}$: clearly N is measurable and integer-valued.

Now consider a behavioral strategy vector σ , and a fixed integer N . Let $\tilde{\sigma}$ coincide with σ during the first N stages, then be equal to $\bar{\sigma} = (\bar{\sigma}_0, \bar{\sigma}_0, \dots)$. Let us first prove (a) and (b) for $\tilde{\sigma}$. $v_1^{\tilde{\sigma}}(\tilde{h}_t) = v_1^{\tilde{\sigma}}(\omega_t) = \bar{w}(\omega_t)$ satisfies (a) for all $t \geq N$. If $v_1^{\tilde{\sigma}}(h, \omega, s, \tilde{\omega})$ exists, is measurable, and is $\geq \underline{w}_1(\tilde{\omega})$ and $\leq \bar{w}_1(\tilde{\omega})$, then $\int v_1^{\tilde{\sigma}}(h, \omega, s, \tilde{\omega}) dp(\tilde{\omega} \mid \omega, s)$ will exist (integrability of $\underline{w}_1(\tilde{\omega})$ and $\bar{w}_1(\tilde{\omega})$) for all (h, ω, s) , and be measurable in (h, ω, s) . Together with the measurability of $u_1(\omega, s)$ this implies that $v_1^{\tilde{\sigma}}(h, \omega, s)$ exists and is measurable. Similarly, $u_1(\omega, s) + v_1^{\tilde{\sigma}}(h, \omega, s, \tilde{\omega})$ will be measurable, and smaller (resp. larger) than the integrable function

$$\max_{s \in S(\omega)} (u_1(\omega, s) + \bar{w}_1(\tilde{\omega})) \left[\text{resp. } \min_{s \in S(\omega)} u_1(\omega, s) + \underline{w}_1(\tilde{\omega}) \right] (\omega \text{ is fixed}).$$

Thus the integral

$$\int [u_1(\omega, s) + v_1^{\tilde{\sigma}}(h, \omega, s, \tilde{\omega})] \, dp(\tilde{\omega} \mid \omega, s) d\tilde{\sigma}(s \mid h, \omega)$$

will be well defined and absolutely convergent; the measurability of the integrand will again imply that the result is measurable in (h, ω) . Thus, by (7.c), $v_1^{\tilde{\sigma}}(h, \omega)$ will exist, be finite-valued, measurable, and equal to this integral. In particular, together with our formula for $v_1^{\tilde{\sigma}}$, this implies (via Fubini) that

$$v_1^{\tilde{\sigma}}(h, \omega) = \int v_1^{\tilde{\sigma}}(h, \omega, s) d\sigma_1(s_1 \mid h, \omega) d\sigma_2(s_2 \mid h, \omega) \dots$$

Thus, to finish proving that if $v_1^{\tilde{\sigma}}(\tilde{h})$ satisfies (a) for $\tilde{h} \in \tilde{H}_{t+1}$, then it satisfies it for $\tilde{h} \in \tilde{H}_t$ and (b) is satisfied for $\tilde{h} \in \tilde{H}_t$ and $h \in H_{t+1}$, there only remains to show that $\underline{w}_1(\omega) \leq v_1^{\tilde{\sigma}}(h, \omega) \leq \bar{w}_1(\omega)$. It is sufficient to prove, e.g., the upper bound. Since $\bar{w}_1(\omega) \geq u_1(\omega, s) + \int \bar{w}_1(\tilde{\omega}) dp(\tilde{\omega} \mid \omega, s) \, \forall s \in$

$S(\omega)$ (cf supra) and $v_1^{\tilde{\sigma}}(h, \omega, s, \tilde{\omega}) \leq \bar{w}_1(\tilde{\omega})$, we have by our formula for $v_1^{\tilde{\sigma}}(h, \omega, s)$ that $v_1^{\tilde{\sigma}}(h, \omega, s) \leq \bar{w}_1(\omega)$. Averaging over s with the strategies yields then $v_1^{\tilde{\sigma}}(h, \omega) \leq \bar{w}_1(\omega)$.

Thus (a) and (b) are fully established for $\tilde{\sigma}$.

Now consider the strategy σ . By (7.b), we have, for $t \geq N$, $\bar{v}_1^{\sigma^{ht}}(\omega) \leq \bar{w}_1(\omega) = \bar{v}_1^{\tilde{\sigma}^{ht}}(\omega)$; i.e., $\bar{v}_1^{\sigma}(h_t, \omega) \leq v_1^{\tilde{\sigma}}(h_t, \omega)$. Since until $t = N$, both σ and $\tilde{\sigma}$ agree, and since (a) and (b) (including b.3) are proved for $\tilde{\sigma}$, (7.c) will imply that the above inequality will hold for $t(< N)$ as soon as it holds for $t+1$. Thus $\bar{v}_1^{\sigma}(\tilde{h}) \leq v_1^{\tilde{\sigma}}(\tilde{h})$ for all \tilde{h} . Similarly, one will have $\underline{v}_1^{\sigma}(\tilde{h}) \geq v_1^{\sigma'}(\tilde{h})$ for all \tilde{h} , where σ' coincides with σ during the first N stages, then with $(\underline{\sigma}_0, \underline{\sigma}_0, \underline{\sigma}_0, \dots)$.

Let us more explicitly subscript $\tilde{\sigma}$ and σ' with N , and show that $v_1^{\tilde{\sigma}^N}(\tilde{h})$ decreases with N (similarly, $v_1^{\sigma'^N}(\tilde{h})$ will increase with N). Clearly, for histories $\tilde{h} = (h, \omega)$ of length N (or larger) $v_1^{\tilde{\sigma}^N}(\tilde{h}) = \bar{w}_1(\omega) \geq v_1^{\tilde{\sigma}^{N+1}}(\tilde{h})$. For shorter histories, $\tilde{\sigma}_N$ and $\tilde{\sigma}_{N+1}$ coincide, so the inequality follows by backward induction, using our formulas of (b) for $v^{\tilde{\sigma}}$.

Now consider $\varphi(\tilde{h}) = v_1^{\tilde{\sigma}^N}(\tilde{h}) - v_1^{\sigma'^N}(\tilde{h})$: those are positive, measurable functions on histories. For \tilde{h} of length N or more, $\varphi(\tilde{h}) = \bar{w}_1(\omega) - w_1(\omega) = w_0(\omega)$ as defined above. For shorter histories, $\tilde{\sigma}_N$ and σ'_N coincide; thus, using our formulas (b)—with $\tilde{h} = (h, \omega) \in \tilde{H}_t, t < N$ —for $v^{\tilde{\sigma}}$ and $v^{\sigma'}$, we get that $\varphi(h, \omega) = \int \varphi(h, \omega, s, \tilde{\omega}) dp(\tilde{\omega} | \omega, s) d\sigma(s | h, \omega)$. Since $w_{i+1}(\omega) = \max_{s \in S(\omega)} \int w_i(\tilde{\omega}) dp(\tilde{\omega} | \omega, s)$, the inequality $\varphi(h, \omega) \leq w_0(\omega)$ for h of length N yields by backward induction $\varphi(h, \omega) \leq w_{N-k}(\omega)$ for h of length k . Since w_i converges to zero, we conclude that, as $N \rightarrow +\infty$, the $v_1^{\tilde{\sigma}^N}(\tilde{h})$ and the $v_1^{\sigma'^N}(\tilde{h})$ converge (monotonically) to the same limit. Since $v_1^{\sigma'^N} \leq \underline{v}_1^{\sigma}(\tilde{h}) \leq \bar{v}_1^{\sigma}(\tilde{h}) \leq v_1^{\tilde{\sigma}^N}(\tilde{h})$, we conclude first that $v_1^{\sigma}(\tilde{h})$ exists, is measurable in \tilde{h} , and satisfies $\underline{v}_1^{\sigma}(\omega) \leq v_1^{\sigma}(h, \omega) \leq \bar{w}_1(\omega)$. Similarly, the recursion formulas (b) for v_1^{σ} now immediately follow by monotone convergence from those for $v_1^{\sigma'^N}$ and $v_1^{\tilde{\sigma}^N}$.

Thus (a) and (b) are established.

For (c), fix the initial state ω , and let $N = N(\omega)$. Since σ and τ coincide during the first N stages, we get that both $v_1^{\sigma}(\omega)$ and $v_1^{\tau}(\omega)$ lie between $v_1^{\sigma'^N}(\omega)$ and $v_1^{\tilde{\sigma}^N}(\omega)$. Since $0 \leq v_1^{\tilde{\sigma}^N}(\omega) - v_1^{\sigma'^N}(\omega) = \varphi(\omega) \leq w_N(\omega) < \epsilon$, we obtain indeed that $|v_1^{\sigma}(\omega) - v_1^{\tau}(\omega)| < \epsilon$. The present function $N(\omega)$ was constructed for player 1, say $N_1(\omega)$. If instead one uses $\max_n N_n(\omega)$, one indeed obtains $\|v^{\sigma}(\omega) - v^{\tau}(\omega)\| < \epsilon$.

3.6. SUBGAME PERFECTION: THE THEOREM

In the previous paragraphs we have shown that, under our assumptions, the payoff function is unambiguously defined, whatever the strategy spaces are, and that Nash equilibria of the game restricted to behavioral strategies are still Nash equilibria whatever strategy spaces are considered, and thus qualified unambiguously as Nash equilibria of the stochastic game. We have also derived the basic properties of the payoff function $v^\sigma(h)$ and $v^\sigma(\tilde{h})$ conditional to finite histories h and \tilde{h} (Propositions 4 and 5).

For subgame perfection, there is another potential source of ambiguity to resolve: since a subgame is a subtree such that any position in (or out of) the subtree is contained in an information set wholly in (or out of) the subtree, and since the information sets of the players are the measurable subsets in the space of histories, this would lead us to require that subgames be measurable. A great number of different possible concepts of subgame perfection would arise, according to the different σ -fields considered on H . Note that points in the space of histories do not have to be measurable: even just in Ω , \mathcal{A} may very well have no atoms at all and may also not separate points. We here take the strongest conceivable definition of subgame perfection, and we will see that in fact there is no ambiguity after all.

Definition 2 A “subgame-perfect equilibrium” is a strategy vector σ such that, for any finite history \tilde{h} , the vector of conditional strategies $\sigma^{\tilde{h}}$ is a Nash equilibrium.

Corollary 3 Given a subgame-perfect equilibrium σ , the strategy vectors σ^h and $\sigma^{\tilde{h}}$ are also subgame-perfect equilibria for any finite histories h and \tilde{h} .

Proposition 6 (a) A strategy vector σ is a subgame-perfect equilibrium if and only if, given any finite history \tilde{h} , no player can profitably deviate from $\sigma^{\tilde{h}}$ by deviating just in the first stage.
 (b) In other words, iff $\forall \tilde{h}$, the $\sigma_n(ds_n \mid \tilde{h})$ form a Nash equilibrium of the game $v^\sigma(\tilde{h}; (s_1, \dots, s_n, \dots))$.

Proof. The equivalence of the two statements follows from Proposition 5.b. The necessity is obvious. Assume that the conditions are satisfied and σ is not a subgame-perfect equilibrium: there exists a history \tilde{h}_0 , such that $\sigma^{\tilde{h}_0}$ is not a Nash equilibrium. Thus one player—say 1—has a profitable deviation, say τ , in the game starting after \tilde{h}_0 .

τ can be extended as follows to a strategy $\tilde{\tau}$ of player 1 in the full game. τ , being a strategy after \tilde{h}_0 , is a (measurable) function on partial histories $h' = (s_0, \omega_1, s_1, \omega_2, s_2, \omega_3, \dots)$, since ω_0 is fixed in this game. For $s_0 \notin S(\omega_0)$, τ is not defined. Define it for those s_0 by $\tau(s_0, \tilde{h}) = \tau(\bar{s}_0, \tilde{h})$, for a fixed \bar{s}_0 in $S(\omega_0)$. Now τ is well defined for any such partial history h' .

Define $\bar{\tau}(\omega, h') = \tau(h')$ for any ω : $\bar{\tau}$ is a well-defined behavioral strategy. Finally, if $\tilde{h}_0 \in \tilde{H}_{t_0}$, let, for any $h \in H_{t_0}$, and any \tilde{h} , $\bar{\tau}(h, \tilde{h}) = \bar{\tau}(\tilde{h})$, and for $\tilde{h} \in \tilde{H}_s$, $s < t_0$, let $\bar{\tau}(\tilde{h}) = \sigma_1(\tilde{h})$: $\bar{\tau}$ is a strategy of player 1 in the full game, which coincides with σ_1 before t_0 , and with τ after \tilde{h}_0 .

By Proposition (5.c), we can change $\bar{\tau}$ so as to coincide with σ_1 after a sufficiently large number N of stages, and it will still be a profitable deviation after \tilde{h}_0 . Then there is a history $\tilde{h}_1 = (h, \omega)$ of maximal length t ($\leq \text{length}(\tilde{h}_0) + N$) extending \tilde{h}_0 such that $\bar{\tau}$ is still a profitable deviation after \tilde{h}_1 . There is no loss in changing $\bar{\tau}$ so as to coincide with σ before stage t . Define $\bar{\tau}$ such as to coincide with $\bar{\tau}$ at stage t , and with σ at all other stages; since $\bar{\tau}$ was not a profitable deviation after any history \tilde{h} of length $> t$, Proposition 5.b implies that $\bar{\tau}$ is a fortiori a profitable deviation after \tilde{h}_1 .

This proves the proposition.

Corollary 4 *If σ is not subgame perfect, then for some t and some $E \in \tilde{\mathcal{H}}_t$ and for some player n , there exists a strategy τ_n that coincides with σ_n on all partial histories except for $\tilde{h} \in E$, and s.t. for each of those, τ_n is a profitable deviation from σ_n .*

Remark. E satisfies the strictest measurability requirements; this shows that, even with the weakest concept of subgame perfection, σ would not be subgame perfect.

Proof. The essential part was done in the above proof, which proved a bit more than strictly necessary. Consider the deviation $\bar{\tau}$ on histories $\tilde{h} \in \tilde{H}_t$: the set $E = \{\tilde{h} \mid v_n^{\bar{\tau}}(\tilde{h}) > v_n^{\sigma}(\tilde{h})\}$ belongs to $\tilde{\mathcal{H}}_t$ (Proposition 5.a.2), and contains \tilde{h}_1 . Define then τ_n to coincide with $\bar{\tau}$ on E , and with σ_n everywhere else. This proves the corollary. (Even if, under requirements (A'), the original σ -field \mathcal{A} had been extended to $\mathcal{B}(\mathcal{A})$ (cf Section 2), E would contain a set \tilde{E} containing \tilde{h}_1 and belonging to the product σ -field generated by the original σ -fields \mathcal{A} (consider the unit mass at \tilde{h}_1 , and that E is universally measurable). Thus the “subgame” can really be chosen to satisfy the strictest measurability requirements.)

We thus see that we also obtain a completely unambiguous concept of subgame-perfect equilibria, which according to the proposition coincide with the “backward induction equilibria.”

Our result is:

Theorem 1 *There exist subgame-perfect equilibria.*

The proof will be given in the next three sections.

4. Separability

We show here that it is sufficient to prove the theorem when the σ -field \mathcal{A} is separable.

Proposition 7 *One can assume \mathcal{A} separable. More precisely:*

- a) *There exists a minimal σ -field $\tilde{\mathcal{F}}$ for which all assumptions of Section 2.2.1 and of Section 3, Proposition 3 are satisfied. $\tilde{\mathcal{F}}$ is separable, and $\tilde{\mathcal{F}} \subseteq \mathcal{A}$ (resp. $\mathcal{B}(\mathcal{A})$) if the original model satisfies requirements (A) (resp. (A)') of Section 2.*
- b) *The functions \underline{w} and \bar{w} of Section 3 are $\tilde{\mathcal{F}}$ -measurable, and equal to the corresponding functions for $(\Omega, \tilde{\mathcal{F}})$.*
- c) *Any subgame-perfect equilibrium for $(\Omega, \tilde{\mathcal{F}})$ is one for the original model.*

Proof. a) Note that, under the assumptions of Section 3.3.2, the σ -field \mathcal{F} of Section 2 fills the bill. Under the more general assumptions of Section 3, Proposition 3, however, nothing guarantees that the v^{σ_0} of Proposition 3.a will be \mathcal{F} -measurable.

Consider thus an arbitrary strategy σ , and the corresponding function v^σ (Proposition 5). v^σ generates, together with \mathcal{F} , a separable σ -field \mathcal{F}' . Enlarge \mathcal{F}' , as in the proof of Section 2, Proposition 1.c, to the minimal σ -field \mathcal{F}_σ containing \mathcal{F}' and such that p is a transition probability for \mathcal{F}_σ : \mathcal{F}_σ is separable, satisfies the assumption of Section 2.2.1, and those of Section 3, Proposition 3. Indeed, 3.a is satisfied by construction, 3.c is satisfied (when keeping all v_σ 's the same, at least for those σ 's that remain strategies) because upper integrals can only increase when σ -fields are decreased, 3.d and 3.e remain true because they were so in the original model, and 3.b because the proof of Section 3, Lemma 8 (which does not depend on 3.b) shows that strategies $\bar{\sigma}_n$ and $\underline{\sigma}_n$ can be chosen to be measurable with respect to any σ -field for which all other assumptions of Section 3, Proposition 3 and Section 2.2.1 hold, in particular \mathcal{F}_σ , and such as to guarantee the same functions \bar{w}_n and \underline{w}_n (cf last paragraph of the proof): since those are, by 3.b, valid bounds in the model with the larger σ -field, they are a fortiori so with the smaller σ -field. This by the way also establishes (b).

Since \mathcal{F}_σ satisfies all assumptions, Section 3, Proposition 5 implies that, for any \mathcal{F} -measurable strategy τ (such strategies exist by Section 2, Corollary 2), v^τ will be \mathcal{F}_σ -measurable, so that we will have $\mathcal{F}_\tau \subseteq \mathcal{F}_\sigma$ by the minimality property of \mathcal{F}_τ . Thus, for any two \mathcal{F} -measurable strategies τ_0 and τ_2 , we will have $\mathcal{F}_{\tau_1} = \mathcal{F}_{\tau_2}$: denote this σ -field by $\tilde{\mathcal{F}}$. $\tilde{\mathcal{F}}$ satisfies all assumptions, is separable, and $\tilde{\mathcal{F}} \subseteq \mathcal{F}_\sigma$ for any strategy σ , hence its minimality. The inclusion of $\tilde{\mathcal{F}}$ in \mathcal{A} (resp. in $\mathcal{B}(\mathcal{A})$) follows from its minimality, since those were σ -fields with respect to which all assumptions held (cf Section 2, Remark 1). This proves (a). (Note that, while we showed that the original functions \bar{w}_n and \underline{w}_n are measurable w.r.t. the minimal $\tilde{\mathcal{F}}$, we could even have used Proposition 5.c to show that $\tilde{\mathcal{F}}$ is in fact generated as the minimal σ -field for which Section 2.2.1 holds and for which those functions are measurable.)

c) follows, by application of Proposition 6.b., because our above proof shows that the assumptions remain true on the smaller σ -field while keeping the same functions v^σ —for those σ 's that remain strategy vectors—hence the same $v^\sigma(h)$.

Thus, we henceforth assume \mathcal{A} separable, in addition to Section 2.2.1 and Section 3, Proposition 3.

5. The Candidate Set

Denote by P the Euclidean space of vector payoffs. We write for short \mathcal{K}_Y^* instead of $\mathcal{K}_Y \setminus \phi$, for any Y . Let Γ —the “space of games”—denote the subspace of $\mathcal{K}_{P \times \bar{S}}^*$ consisting of the graphs of continuous P -valued functions defined on sets $\Pi_n C_n$, with $C_n \in \mathcal{K}_{\bar{S}_n}^*$.

The mapping Φ of Section 3, Proposition 4 is such that, for each ω and f , $[\Phi(f)](\omega)$ is a continuous, P -valued function on $S(\omega)$. Any such function can be equivalently described by its graph, which belongs to Γ : redefine ϕ then as having values in Γ . Since all conclusions of Section 3, Proposition 4 are for fixed ω , and since on the space of continuous functions on a fixed set $S(\omega)$ the Hausdorff topology on the graphs coincides with the uniform topology, the conclusions of Section 3, Proposition 4 are not affected by this redefinition.

Finally, define for Borel maps $C: (\Omega, \mathcal{A}) \rightarrow \mathcal{K}_P^*$ satisfying $C(\omega) \subseteq W(\omega)$ ($C \in \text{Dom}(\Phi)$ for short), $[\Phi(C)](\omega)$ to be the closure in Γ of $\{[\Phi(f)](\omega) \mid f \text{ is a measurable selection from } C\}$. We obtain then

Lemma 9 a) Γ is a Polish space.

b) $C \in \text{Dom}(\Phi) \Rightarrow \Phi(C)$ is a measurable map from (Ω, \mathcal{A}) to \mathcal{K}_Γ^* .

c) $C_i \in \text{Dom}(\Phi), C_i \rightarrow C$ pointwise $\Rightarrow \Phi(C_i) \rightarrow \Phi(C)$ pointwise.

Proof. a) $\mathcal{K}_{P \times \bar{S}}^*$ is Polish by [5], Proposition 5.a. Let $F_k = \{K \in \mathcal{K}_{P \times \bar{S}}^* \mid \exists p_1, p_2, s: (p_1, s) \in K, (p_2, s) \in K, \|p_1\| \leq k, \|p_2\| \leq k, \|p_1 - p_2\| \geq k^{-1}\}$: clearly, F_k is closed, and $\cup F_k$ is the complement of the set of graphs of continuous functions from a compact subset of \bar{S} to P . Thus the latter set is a G_δ in $\mathcal{K}_{P \times \bar{S}}^*$, hence Polish too. Finally, Γ is closed in the latter set, because the projection to \bar{S} is continuous and $\Pi_n \mathcal{K}_{\bar{S}_n}^*$ is a closed subspace of $\mathcal{K}_{\bar{S}}^*$ ([5], Proposition 6.d and 6.f). Hence Γ is Polish too. The \mathcal{G} -measurability of $\Phi(f)$ —as a map from G to P —follows immediately from the assumptions, and its continuity in s from Section 3, Proposition 4.a. Section 2, Corollary 1 then implies directly that $\Phi(f)$ is a Borel map to Γ for any $f \in \mathcal{F}$. Hence, by Section 3, Proposition 4, (a) and (b), all assumptions of [5], Proposition 8 are satisfied: (a) and (c) of that proposition yield then (b) and (c) of Lemma 9. This finishes the proof.

From Γ , we have natural projections to each set $\mathcal{K}_{\bar{S}_n}^*$; denote the image of $\gamma \in \Gamma$ by $S_n(\gamma)$. Similarly, $S(\gamma) = \Pi_n S_n(\gamma)$ is the projection on \bar{S} . Hence ([5], Proposition 6d) $S_n(\gamma)$ and $S(\gamma)$ are continuous functions on Γ . Denote also by $v(\gamma)$ or v^γ the continuous P -valued function on $S(\gamma)$ that corresponds to γ .

Denote by Σ_n the space of probabilities on \bar{S}_n , with the weak*-topology. Let $\Sigma = \Pi_n \Sigma_n$, $X = \{(\gamma, \sigma) \in \Gamma \times \Sigma \mid \sigma_n \text{ is carried by } S_n(\gamma)\}$, $E = \{(\gamma, \sigma) \in X \mid \sigma \text{ is a Nash equilibrium of } \gamma\}$, and $\pi : X \rightarrow P : \pi(\gamma, \sigma) = \int v^\gamma(s) d\sigma(s)$. Finally, let $\hat{X}(\gamma) = \{\gamma\} \times \{\sigma \mid (\gamma, \sigma) \in X\}$, $\hat{E}(\gamma) = \{\gamma\} \times \{\sigma \mid (\gamma, \sigma) \in E\}$.

Lemma 10 (*uppersemicontinuity of the equilibrium correspondence*)

- a) E and X are closed in $\Gamma \times \Sigma$, hence Polish.
- b) \hat{X} and \hat{E} are u.s.c. maps from Γ to $\mathcal{K}^*(X)$ and $\mathcal{K}^*(E)$ respectively.
- c) π is continuous.

Proof. Note that the Σ_n 's are compact metric; hence it is sufficient to check closedness and continuity along convergent sequences. A direct proof could be given, but it is more convenient to use Skohorod's theorem: if $\sigma^i \in \Sigma_n$ converges to σ^∞ , there exists a sequence of \bar{S}_n -valued random variables $s^i(t)$ on $([0, 1], \text{Lebesgue measure})$ such that $s^i(t)$ converges a.e. to $s^\infty(t)$ and such that σ^i (resp. σ^∞) is the distribution of $s^i(t)$ (resp. $s^\infty(t)$). Thus, if $\sigma^i \in \Sigma$ converges to σ^∞ , we get such random variables $s_n^i(t_n)$ for each player n , where the t_n 's are independent uniform random variables. Since any Borel subset of $[0, 1]$ with Lebesgue measure 1 has a Borel isomorphism with $[0, 1]$ preserving Lebesgue measure, there is no loss in further assuming the $s^i(t)$ to be Borel measurable, with values in the support of σ^i , and converging everywhere to $s^\infty(t)$.

Hence, if $(\gamma^i, \sigma^i) \in X$ converges to $(\gamma^\infty, \sigma^\infty)$ in $\Gamma \times \Sigma$, we have seen that the $S_n(\gamma^i)$'s converge (Hausdorff) to $S_n(\gamma^\infty)$. Since for each t , $s_n^i(t) \in S_n(\gamma^i)$ and converges to $s_n^\infty(t)$, it follows that $s_n^\infty(t) \in S_n(\gamma^\infty)$; hence σ_n^∞ is carried by $S_n(\gamma^\infty)$: $(\gamma^\infty, \sigma^\infty)$ belongs indeed to X ; hence the closedness of X is proved. Further, $\pi(\gamma^i, \sigma^i) = E[v^{\gamma^i}(s_1^i(t_1), s_2^i(t_2), \dots)] = Ev^{\gamma^i}(s^i(t))$ in vector notation. Since $(v^{\gamma^i}(s^i(t)), s^i(t)) \in \gamma^i$, and since $\gamma^i \rightarrow \gamma^\infty$, the sequence is compact, and any limit point belongs to γ^∞ . But the second coordinate converges to $s^\infty(t)$, so there is only one possible value for the limit points in γ^∞ , i.e., $(v^{\gamma^\infty}(s^\infty(t)), s^\infty(t))$. Thus $v^{\gamma^i}(s^i(t))$ converges to $v^{\gamma^\infty}(s^\infty(t))$. Finally, $\gamma^i \rightarrow \gamma^\infty$ implies that $\gamma^\infty \cup (\cup_i \gamma^i)$ is compact; hence the sequence $v^{\gamma^i}(s^i(t))$ is uniformly bounded: the dominated convergence theorem implies therefore that $Ev^{\gamma^i}(s^i(t))$ converges to $Ev^{\gamma^\infty}(s^\infty(t))$, i.e., $\pi(\gamma^i, \sigma^i) \rightarrow \pi(\gamma^\infty, \sigma^\infty)$: continuity of π is also established.

Closedness of E now follows also: assume $(\gamma^i, \sigma^i) \in E$, and $(\gamma^\infty, \sigma^\infty) \notin E$. Then one player, say 1, has a profitable deviation: $\exists s^\infty \in S_1(\gamma^\infty)$ s.t.

$\pi_1(\gamma^\infty; s^\infty, \sigma_2^\infty, \sigma_3^\infty, \dots) > \pi_1(\gamma^\infty; \sigma_1^\infty, \sigma_2^\infty, \dots)$. Since $S_1(\gamma^i) \rightarrow S_1(\gamma^\infty)$, $\exists s^i \in S_1(\gamma^i)$ such that $s^i \rightarrow s^\infty$. But then, by the continuity of π , we would have $\pi_1(\gamma^i; s^i, \sigma_2^i, \sigma_3^i, \dots) > \pi_1(\gamma^i; \sigma_1^i, \sigma_2^i, \dots)$, a contradiction.

Since Γ is Polish (Lemma 9), $\Gamma \times \Sigma$ is too, as a product of two Polish spaces; hence X and E are Polish, as closed (hence G_δ) subsets of a Polish space. \hat{X} and \hat{E} have nonempty values because any such game has an equilibrium point (we have just shown in particular that the payoff function is jointly continuous (and is clearly multilinear) on the product of the compact mixed strategy spaces), and uppersemicontinuity is an immediate consequence of the compactness of Σ and of the closedness of X and E .

\hat{X} and \hat{E} will also denote the induced maps on \mathcal{K}_Γ^* , and similarly π will still denote the induced map on \mathcal{K}_X^* or its restriction to the subspace \mathcal{K}_E^* :

Corollary 5 a) π is a continuous map from \mathcal{K}_X^* (or \mathcal{K}_E^*) to \mathcal{K}_P^* .

b) \hat{X} and \hat{E} are u.s.c. Borel maps from \mathcal{K}_Γ^* to \mathcal{K}_X^* and \mathcal{K}_E^* resp.

c) \mathcal{K}_Γ^* , \mathcal{K}_X^* , \mathcal{K}_E^* and \mathcal{K}_P^* are Polish.

Proof. (c) follows from Lemma 10.a, Lemma 9.a and [5], Proposition 5.a.

(a) follows from Lemma 10.c and from [5], Proposition 6.d. (b) follows from Lemma 10.b (+ (a) for Polishness of X and E) and from [5], Proposition 5, (b) and (c).

Proposition 8 *There exists a measurable map N from (Ω, \mathcal{A}) to \mathcal{K}_P^* , which is in the domain of Φ , such that*

a) $N = \pi \circ \hat{E} \circ \Phi(N)$.

b) $\forall \epsilon > 0$, there exists a Borel strategy vector σ in the game where the initial state is in the graph of N instead of just in Ω —thus $\sigma(\tilde{h})$ where the \tilde{h} 's are of the form $(p, g_0, g_1, \dots, g_t, \omega)$ with $p \in N(\omega_0)$, such that $\|v^\sigma(p, \omega_0) - p\| \leq \epsilon$ uniformly on the graph of N .

Proof. First note that, by definition of Φ and of the function $\underline{w}_n, \bar{w}_n$, any measurable map N to \mathcal{K}_P^* satisfying (b) lies in the domain of Φ . We first construct a maximal such map. Fix a stationary strategy $\sigma = (\sigma_0, \sigma_0, \dots)$; v^σ is measurable; hence $N_0(\omega) = \{v^\sigma(\omega)\}$ satisfies our requirements. Given N_{k-1} satisfying our requirements, let $X_k(\omega) = \hat{X}[(\Phi(N_{k-1}))(\omega)]$, and $N_k(\omega) = \pi(X_k(\omega))$. By Lemma 9, and by the above corollary, X_k and N_k are measurable maps to \mathcal{K}_X^* and \mathcal{K}_P^* resp.; further $N_0(\omega) \subseteq N_1(\omega)$, since the choice $\sigma_0 \in \Sigma$ is available and would yield $v^\sigma(\omega)$. Hence, by induction, $N_{k-1}(\omega) \subseteq N_k(\omega)$. Finally, N_k satisfies (b): let σ_{k-1} satisfy (b) with $\epsilon/2$ on the graph of N_{k-1} . Let f_i be a sequence of Borel selections from the graph of N_{k-1} such that the $[\Phi(f_i)](\omega)$ are dense in $[\Phi(N_{k-1})](\omega)$ for each ω ([5], Proposition 8.b). Also, let φ be a Borel map from the graph of N_k

to X such that everywhere $\varphi(p, \omega) \in X_k(\omega)$, $p = \pi(\varphi(p, \omega))$ ([5], Proposition 9). Denote the coordinates of φ in Σ and Γ resp. by $\sigma_0(p, \omega)$ and $\gamma(p, \omega)$: those are Borel functions on the graph of N_k such that $\sigma_0(p, \omega)$ is everywhere a strategy vector in $\gamma(p, \omega)$ with payoff p , and such that $\gamma(p, \omega) \in (\Phi(N_{k-1}))(\omega)$. Let $g_i = \Phi(f_i) : \sigma_0(p, \omega)$ is still a strategy vector in all $g_i(\omega)$, with $p_i(p, \omega) = \pi(g_i(\omega), \sigma_0(p, \omega))$ having $p = \pi(\gamma(p, \omega), \sigma_0(p, \omega))$ as limit point (Lemma 10.c). Let $T(p, \omega) = \min\{i \mid \|p_i(p, \omega) - p\| \leq \epsilon/2\}$: $T(p, \omega)$ is Borel, so $\sigma_k(p) = (\sigma_0(p, \omega_0), \sigma_{k-1}[f_{T(p, \omega_0)}(\omega_1), \omega_1])$ is a Borel strategy for the game having the graph of N_k as a set of initial states; since $\|v^{\sigma_{k-1}(f_{T(p, \omega_0)})} - f_{T(p, \omega_0)}\| \leq \epsilon/2$ and since $\|p_{T(p, \omega)}(p, \omega) - p\| \leq \epsilon/2$, the recursion formula (Section 3, Proposition 5.b) implies $\|v^{\sigma_k(p)} - p\| \leq \epsilon$. Thus the N_k form an increasing sequence of correspondences satisfying our requirements. In particular, it follows that $N_k(\omega) \subseteq W(\omega)$. Since an increasing sequence of compact sets, contained in a fixed compact set, converges in the Hausdorff topology, it follows that N_k converges pointwise to a Borel map, say N_∞ from (Ω, \mathcal{A}) to \mathcal{K}_P^* . Since the graphs G_k of N_k are Borel ([5], Proposition 9.b), we obtain a Borel strategy σ defined on $\cup_k G_k = G_\infty$ by $\sigma(p) = \sigma_k(p)$ on $G_k \setminus G_{k-1}$ ($G_{-1} = \emptyset$); then for each $(p, \omega) \in G_\infty$, $\|v^{\sigma(p)}(\omega) - p\| \leq \epsilon$. We now obtain a similar strategy σ_∞ on the graph of N_∞ , which is for each ω the closure of $G_\infty(\omega)$ (Lemma 10.c). Rank in one sequence all functions f_i we have met, at all stages of the induction. Let $T(p, \omega) = \min\{i \mid \|f_i(\omega) - p\| \leq \epsilon\}$, which is Borel on the graph of N_∞ , and let $\sigma_\infty(p, \omega) = \sigma(f_{T(p, \omega)}(\omega), \omega)$. Thus N_∞ satisfies our requirements. By Lemma 9, $\Phi(N_k)$ converges pointwise to $\Phi(N_\infty)$; hence trivially $X_k(\omega)$ converges pointwise to $X_\infty(\omega) = \hat{X}[(\Phi(N_\infty))(\omega)]$. Thus, by Corollary 5(a), $N_k = \pi \circ X_k$ converges pointwise to $\pi \circ X_\infty$; $N_\infty(\omega) = \pi(X_\infty(\omega))$.

Let us now forget this whole sequence, and write N_0 for N_∞ : we have a Borel map N_0 from (Ω, \mathcal{A}) to \mathcal{K}_P^* that satisfies (b), and is such that $N_0 = \pi \circ \hat{X} \circ [\Phi(N_0)]$. Hence $N_1 = \pi \circ \hat{E} \circ [\Phi(N_0)] \subseteq N_0$: defining inductively $E_k = \hat{E} \circ [\Phi(N_k)]$, $N_{k+1} = \pi \circ E_k$ we find that the E_k and the N_k are decreasing sequences of Borel maps (Lemma 9 and Corollary 5) from (Ω, \mathcal{A}) to \mathcal{K}_E^* and to \mathcal{K}_P^* respectively. Hence they are pointwise convergent, say to E_∞ and N_∞ , which are thus also Borel maps from (Ω, \mathcal{A}) to \mathcal{K}_E^* and \mathcal{K}_P^* . By Lemma 9, $\Phi(N_k)$ converges pointwise to $\Phi(N_\infty)$. Hence, by Corollary 5.b, for fixed ω , any open set U of E containing $\hat{E}([\Phi(N_\infty)](\omega))$ contains all $E_k(\omega) = \hat{E}([\Phi(N_k)](\omega))$, for k sufficiently large. Since also $N_\infty \subseteq N_k$ implies that $\Phi(N_\infty) \subseteq \Phi(N_k)$ and hence that $\hat{E}([\Phi(N_\infty)](\omega)) \subseteq E_k(\omega)$, we find that $E_\infty = \hat{E} \circ [\Phi(N_\infty)]$. Finally, $E_k \rightarrow E_\infty$ implies, by Corollary 5.a, that $N_{k+1} = \pi \circ E_k$ converges pointwise to $\pi \circ E_\infty$, hence $N_\infty = \pi \circ E_\infty$. Thus N_∞ is a Borel map from (Ω, \mathcal{A}) to \mathcal{K}_P^* , satisfying $N_\infty = \pi \circ \hat{E} \circ \Phi(N_\infty)$, and also satisfying (b) since $N_\infty \subseteq N_0$ which satisfies (b). This proves

Proposition 8.

6. The Equilibrium Strategies

Given the map N of Section 5, Proposition 8, denote by \mathcal{N} the set of Borel selections from N . Let also, for $g \in G$, $\Psi(g) = \{\int f(\omega)p(d\omega \mid g) \mid f \in \mathcal{N}\}$, and denote the graphs of $N \times S[\omega \rightarrow N(\omega) \times S(\omega)]$, N and Ψ resp. by \mathcal{H} , \mathcal{P} and \mathcal{F} .

- Lemma 11** a) Ψ is a Borel map from (G, \mathcal{G}) to \mathcal{K}_P^* , and \mathcal{F} is Borel in $\Omega \times \bar{S} \times P$;
 b) \mathcal{P} is a Borel in $\Omega \times P$;
 c) \mathcal{H} is a Borel in $\Omega \times P \times \bar{S}$;
 d) there exist Borel maps $\sigma_n(p, \omega)$ from \mathcal{P} to Σ_n and $\psi(p, g)$ from \mathcal{H} to P such that $\psi(p, g) \in \Psi(g)$ and the $\sigma_n(p, \omega)$ from a Nash equilibrium with payoff p of the game with pure strategy sets $S_n(\omega)$ and a (continuous) payoff function $u(\omega, s) + \psi(p, \omega, s)(s \in S(\omega))$;
 e) there exists a Borel map φ from $\mathcal{F} \times \Omega$ to P such that $\varphi(g, p, \omega) \in N(\omega)$ and $\int \varphi(g, p, \omega)p(d\omega \mid g) = p$.

Proof. (a) follows from [5], Theorem, a.2. and [5], Proposition 9.b.

(b) follows from [5], Proposition 9.b.

(c) The map $\omega \rightarrow (N(\omega), S(\omega)) \in \mathcal{K}_P^* \times \mathcal{K}_{\bar{S}}^*$ is measurable, hence ([5], Prop 6.f) so is the map $\omega \rightarrow N(\omega) \times S(\omega) \in \mathcal{K}_{P \times \bar{S}}^*$. Hence ([5], Proposition 9.b) its graph \mathcal{H} is measurable.

(d) By [5], Proposition 9.c, there exist Borel maps $\sigma_n(p, \omega)$ and $\gamma(p, \omega)$ from \mathcal{P} to Σ_n and to Γ such that the $\sigma_n(p, \omega)$ form a Nash equilibrium with payoff p of $\gamma(p, \omega)$, and $\gamma(p, \omega) \in [\Phi(N)](\omega)$, because $N = \pi \circ \hat{E} \circ \Phi(N)$, and π is continuous (Section 5, Proposition 8.a and Corollary 5.a). For $(p, \omega, s) \in \mathcal{H}$, let $\psi(p, \omega, s) = [\gamma(p, \omega)](s) - u(\omega, s)$, where $[\gamma(p, \omega)](s)$ denotes the value at $s \in S(\omega)$ of the payoff function of the game $\gamma(p, \omega)$. Let \bar{P} denote the one-point compactification of P . Since Γ is a subspace of $\mathcal{K}_{\bar{P} \times \bar{S}}^*$, by definition and by [5], Proposition 6, g, the map $(\gamma, s) \rightarrow \gamma \cap (\bar{P} \times \{s\})$ is Borel from $\Gamma \times \bar{S}$ to $\mathcal{K}_{\bar{P} \times \bar{S}}^*$, by [5], Proposition 6.e. By composition, $(p, \omega, s) \rightarrow (\gamma(p, \omega)) \cap (\bar{P} \times \{s\})$ is Borel on $\mathcal{P} \times \bar{S}$. Since the values of the restriction of this map to \mathcal{H} are the singletons $\{[\gamma(p, \omega)](s)\}$, it follows that $(p, \omega, s) \rightarrow [\gamma(p, \omega)](s)$ is Borel on \mathcal{H} ([5], Proposition 6.a). Since also $u(\omega, s)$ is measurable, we obtain that ψ is a Borel map from \mathcal{H} to P . Clearly, the fact that $\sigma_n(p, \omega)$ is a Nash equilibrium of $\gamma(p, \omega)$ means that $\sigma_n(p, \omega)$ is a Nash equilibrium of the game with pure strategy spaces $S_n(\omega)$ and (continuous) payoff function $u(\omega, s) + \psi(p, \omega, s)$. Finally, $\psi(p, g) \in \Psi(g)$ follows from the compactness of $\Psi(p)$ (point (a)), and from $\gamma(p, \omega) \in [\Phi(N)](\omega)$, which implies that $\psi(p, g)$ is in the closure of $\Psi(g)$.

Finally, point (e) follows from point (c) of the theorem in [5].

Lemma 12 *The graph of N is the union of the graphs of its Borel selections. For any Borel selection $p(\omega)$ from $N(\omega)$, there exist a Borel map p from (H, \mathcal{H}) to P , and a vector of Borel strategies τ_n , such that*

- a) for $\tilde{h} = (\omega) \in \tilde{H}_0, p(\tilde{h}) = p(\omega)$
- b) the $\tau_n(h, \omega)$ form a Nash equilibrium with payoff $p(h, \omega)$ of the game with pure strategy sets $S_n(\omega)$ and continuous payoff function $p(h, \omega, s)$
- c) $p(h, g) = u(g) + \int p(h, g, \omega) dp(\omega \mid g)$
- d) $p(h, \omega) \in N(\omega)$

Proof. The first statement follows from Section 5, Proposition 8 and [5], Proposition 7.c. (The measurability of a function on (Ω, \mathcal{A}) is preserved when changing its value on some given atom (separability).) Now define inductively a Borel function p on (H, \mathcal{H}) , as follows: for $\tilde{h} = (\omega) \in \tilde{H}_0$, let $p(\tilde{h}) = p(\omega)$. Define p on H_{t+1} by $p(h, \omega_t, s_t) = \psi[p(h, \omega_t), (\omega_t, s_t)]$ for $h \in H_t$ and on \tilde{H}_{t+1} by $p(h, g_t, \omega_{t+1}) = \varphi[g_t, p(h, g_t), \omega_{t+1}]$. Lemma 11 immediately implies by induction (composition of Borel functions) that p is everywhere well defined and is Borel, and that $p(h, \omega) \in N(\omega)$. Now redefine $p(h, g)$ by adding $u(g)$ to it.

Now let $\tau_n(h, \omega) = \sigma_n(p(h, \omega), \omega)$: again the measurability of σ_n (Lemma 11.d) and of p imply by composition that τ_n is Borel measurable—it is well defined because $p(h, \omega) \in N(\omega)$. Thus τ_n is a Borel map from histories \tilde{h} to Σ_n . The rest of the statement follows now from Lemma 11.d for (b), and Lemma 11.e for (c).

Since $N(\omega) \neq \emptyset$ (Section 5, Proposition 8), Lemma 12 implies that, in order to finish the proof of the main theorem, it will be sufficient to prove:

Lemma 13 *The strategy vector $\tau = (\tau_n)$ of Lemma 12 is a subgame-perfect equilibrium with $v^\tau(h) = p(h)$ and $v^\tau(\tilde{h}) = p(\tilde{h})$ for any finite histories h and \tilde{h} .*

Proof. Consider the strategy vector τ^k consisting of playing τ for the first k stages, and at all later stages $k + t$ ($t = 0, 1, \dots$), $\tau^k(h, \tilde{h}) = \sigma(p(h, \omega), \tilde{h})$ where $h \in H_k, \tilde{h} \in \tilde{H}_t, \omega$ denotes the first element of \tilde{h} , and σ is the strategy described in Section 5, Proposition 8.b, with $\epsilon = k^{-1}$. This is well defined since $p(h, \omega) \in N(\omega)$ (Lemma 12.d), and is (by composition) a Borel strategy vector. By definition (Section 3, Proposition 3.c) we have, for all n and all $\tilde{h} \in \tilde{H}_k, |v_n^{\tau^k}(\tilde{h}) - p_n(\tilde{h})| \leq k^{-1}$. By the recursion formula of Section 3, Proposition 5.b for v^{τ^k} and by Lemma 12, (b) and (c) for the function p , we obtain now by backward induction that $|v_n^{\tau^k}(\tilde{h}) - p_n(\tilde{h})| \leq k^{-1}$ and $|v_n^{\tau^k}(h) - p_n(h)| \leq k^{-1}$ for all n and all $\tilde{h} \in \tilde{H}_t, h \in H_t$, and all $t \leq k$. By Section 3, Proposition 5.c, for any $\tilde{h}, v^{\tau^k}(\tilde{h})$ converges to $v^\tau(\tilde{h})$ —the $v^\tau(h)$

and the $v^\tau(\tilde{h})$ exist by Section 3, Proposition 5.a.1 and 5.a.2. By points (a.3) and (b.1) of that proposition, this implies by the dominated convergence theorem that also $v^{\tau^k}(h)$ converges to $v^\tau(h)$. Hence $v^\tau(h) = p(h)$ and $v^\tau(\tilde{h}) = p(\tilde{h})$.

Hence the conclusion, by Lemma 12.b and Section 3, Proposition 6.b.

References

1. Arrow, K.J. and Debreu, G. (1954) Existence of an equilibrium for a competitive economy, *Econometrica* **22**, 265–290.
2. Aumann, R.J. (1964) Mixed and behavior strategies in infinite extensive games, in M. Dresher, L.S. Shapley and A.W. Tucker (eds.), *Advances in Game Theory*, Annals of Mathematics Studies 52, Princeton University Press, Princeton, NJ, pp. 627–650.
3. Forges, F. (1986) An approach to communication equilibria, *Econometrica* **54**, 1375–1385.
4. Mas-Colell, A. (1982) The Cournotian foundations of Walrasian equilibrium theory: An exposition of recent theory, in W. Hildenbrand (ed.), *Advances in Economic Theory*, Econometric Society Monographs in Quantitative Economics, Cambridge University Press, Cambridge, pp. 183–224.
5. Mertens, J.-F. (1987) A measurable “measurable choice” theorem, CORE Discussion Paper 8749, Université Catholique de Louvain, Louvain-la-Neuve, Belgium (Chapter 9 in this volume).
6. Nash, J. (1951) Noncooperative games, *Annals of Mathematics* **54**, 286–295.