More on maximum likelihood equilibria for games with random payoffs and participation

Erik J. Balder, Mathematical Institute, University of Utrecht, P.O. Box 80.010, 3508 TA Utrecht, the Netherlands*

November, 2000

Abstract. Maximum likelihood Nash equilibria were introduced by Borm *et al.* (1995) for games with finitely many players and random payoffs. In Voorneveld (1999) random participation was added to the model. These existence results were extended by Balder (2000c) to continuum games with random payoffs and participation. However, in that paper the complicated measurability issue for the central equilibrium likelihood notion was bypassed by using inner probabilities for the central equilibrium likelihood notion. Here those measurability questions are shown to have a quite satisfactory resolution; this makes the maximum likelihood equilibrium notion more natural. Our main results are not only more general than those found in Borm *et al.* (1995) and Voorneveld (1999), but also improve upon them.

Key words. Nash equilibrium, maximum likelihood equilibrium, random payoffs, random participation.

1 Introduction

Consider the following random game that is a variation of the classical game of matching pennies. There are two players, each of whom can choose H (heads) or T (tails) at the start of the game. Only after they have made their choice, a certain stochastic outcome $\omega \in \mathbb{R}_+$, governed by a probability P, is registered. Based on this stochastic outcome, the payoffs are as follows: (H, H) gives payoffs $(1 - \omega, -1)$ [i.e., player 1 receives $1 - \omega$ and player 2 gets -1], (H, T) gives $(-1 + 2\omega, 1 - \omega)$, (T, H) gives $(-1 + \omega, 1)$ and (T, T) gives payoffs $(1, -1 + 3\omega)$. It is easily verified that the following obtains for the four possible pure action profiles:

(i) (H, H) and (T, H) are not Nash equilibria, regardless of the outcome,

(ii) (H, T) is a Nash equilibrium if and only if $\omega \in [1, 2]$,

(iii) (T, T) is a Nash equilibrium if and only if $\omega \in [\frac{2}{3}, 1]$.

*tel: +31-30-2531458, fax: +31-30-2518394, email: balder@math.ruu.nl

Thus, the *equilibrium likelihood* of the action profile (H, T), i.e., the probability of (H, T) being a Nash equilibrium, is the probability that the stochastic outcome lies in $[1, 2]$. Likewise, the equilibrium likelihood of (T, T) is the probability of that outcome ending up in $[\frac{2}{3}, 1]$ under P. If the former probability is larger than the latter, then (H, T) is said to be the *maximum likelihood Nash equilibrium*, and if the latter exceeds the former then this title goes of course to the action profile (T, T). Observe that when P places all its probability mass on the interval $[0, \frac{2}{3})$, then all four action profiles are maximum likelihood Nash equilibria, albeit with equilibrium likelihood zero! Of course, this is indicative of the situation for the classical game of matching pennies, which has no Nash equilibrium in pure action profiles. Notice that the game of matching pennies amounts to choosing P above to be the Dirac probability measure concentrated at 0.

Maximum likelihood Nash equilibria (MLNE for short) for games with random payoffs were first considered by Borm, Cao and García-Jurado in Borm *et al.* (1995). Their measurability and existence results were improved by Voorneveld (1999), who added random participation to their model. In these two references both measurability and existence questions for MLNE are considered for games with finitely many players. Observe that even for a game with just one player the existence question is already relevant; cf. Example 2.5. In Balder (2000c) the existence of MLNE was studied for a quite general continuum game, i.e., a game with a measure space of players in the spirit of Aumann (1964) and Schmeidler (1973), that has random payoffs and participation. This is a natural continuation of the models studied in Balder (1999b, 2001). However, random participation causes certain measurability questions to surface in connection with the central equilibrium likelihood notion. To circumvent these, the device of an inner probability was proposed by Balder (2000c). From a modeling point of view this seems rather artificial. In the present work it is shown that under broad measurability conditions, which also improve those for the much simpler games of Borm *et al.* (1995) and Voorneveld (1999), these measurability questions can be addressed adequately. This justifies the equilibrium likelihood notion and shows that it has a natural form and interpretation.

The setup of this paper is as follows. In section 2 we introduce a random continuum game Γ, define the notion of MLNE and present Theorem 2.3, our main measurability result. This is followed by Theorem 2.8. Its proof given here is quite different from the one in Balder (2000c) and is more in line with the development of a general approach to equilibrium existence given in Balder (1995, 1999a, 1999b, 2001).

2 MLNE for a random continuum game

We shall define a random continuum game Γ with realizations Γ_ω, $\omega \in \Omega$. Here $(\Omega, \mathcal{F}, P)$ is an abstract probability space, the space of stochastic *outcomes*. Let $(T, \mathcal{T}, \mu)$ be a complete and finite measure space; T is the set of all possible *players*. This is entirely in the spirit of the work by Aumann (1964) and Schmeidler

(1973), who introduced (nonatomic) measure spaces of players to model perfect competition. Recall from section 3.5.1 in Dudley (1989) that T can always be split into a purely atomic part T^{pa} and a nonatomic part T^{na}. Here T^{pa} is an at most countable union of sets A_j, with j in some at most countable index set J. Each A_j is a non-null atom of the measure space $(T, \mathcal{T}, \mu)$. If there are only finitely many players in T (as is the case in Borm *et al.* (1995) and Voorneveld (1999)), then one can take $T^{pa} = T$, $T^{na} = \emptyset$ and for $\mathcal{T}$ one takes the power set 2^T (i.e., the collection of all subsets of T). For μ one then uses simply the counting measure on T (i.e., $\mu(A) :=$ number of players in $A \subset T$). To every outcome $\omega \in \Omega$ there corresponds a set $T_\omega \subset T$, which is the set of players *participating* in the game Γ_ω. As the first of several measurability conditions, we suppose the following:

$$D := \{(\omega, t) \in \Omega \times T : t \in T_\omega\} \text{ is } \mathcal{F} \otimes \mathcal{T}\text{-measurable.} \tag{2.1}$$

This ensures the proper randomness of the set of participating players. In particular, by Proposition III.1.2 of Neveu (1965) this implies that T_ω is $\mathcal{T}$-measurable for every $\omega \in \Omega$. Observe that by assumption (2.1) the trace σ-algebra $\mathcal{D} := D \cap (\mathcal{F} \otimes \mathcal{T})$ can alternatively be described as

$$\mathcal{D} = \{C \in \mathcal{F} \otimes \mathcal{T} : C \subset D\}. \tag{2.2}$$

As the *action universe* we use a Hausdorff locally convex topological vector space S that is a Suslin space for its topology (see Dellacherie and Meyer (1975) or Schwartz (1975)). As is well known (e.g., see p. 25 of Balder (2000a) or p. 214 of Balder (1999b)), such a space S can be endowed with a *weak metric* ρ, which is not stronger than the original topology. The very general notion of a Suslin locally convex space includes all separable Fréchet spaces, including all separable Banach spaces, also when they are considered with their weak topology. Of course, for many game theoretical purposes it suffices to take S Euclidean, but even such a drastic simplification does little to reduce the difficulty of the central measurability and topological questions for the continuum games studied here.

As usual, the Borel σ-algebra on S is denoted by $\mathcal{B}(S)$. Let $\Sigma : T \to 2^S$ be a multifunction that specifies for each player $t \in T$ his/her *feasible* action set $\Sigma(t)$. We suppose that Σ satisfies the following measurability condition:

$$\text{gph } \Sigma := \{(t, s) \in T \times S : s \in \Sigma(t)\} \text{ is } \mathcal{T} \otimes \mathcal{B}(S)\text{-measurable,} \tag{2.3}$$

Let E be the set of all $(\omega, t, s) \in \Omega \times T \times S$ such that $t \in T_\omega$ and $s \in \Sigma(t)$. By (2.1) and (2.3), E belongs to $\mathcal{F} \otimes \mathcal{T} \otimes \mathcal{B}(S)$, for it is easy to see that $E = (D \times S) \cap (\Omega \times \text{gph } \Sigma)$. Consequently, the analogue of (2.2) also holds for the trace σ-algebra $E \cap (\mathcal{F} \otimes \mathcal{T} \otimes \mathcal{B}(S))$.

Let $\mathcal{S}_\Sigma$ be the set of all measurable functions $f : T \to S$ such that $f(t) \in \Sigma(t)$ for every t in T. So $\mathcal{S}_\Sigma$ consists precisely of all the measurable selections of the multifunction Σ. These constitute the *a priori action profiles* of the game Γ, where *a priori* refers to the fact that players are supposed to choose their actions *before* observing the stochastic outcome ω, i.e., before knowing which game Γ_ω

is to be played. Such profiles can be seen as complete descriptions of how the players could or should act *a priori*. For $\omega \in \Omega$ we also define $\Sigma_\omega : T_\omega \to 2^S$ to be the restriction $\Sigma_\omega := \Sigma |_{T_\omega}$ of Σ to T_ω. Thus, $\mathcal{S}_{\Sigma_\omega}$ is the set of all measurable selections of the multifunction Σ_ω. At the same time, it is the set of all *action profiles* for the game Γ_ω, which is only played by the players in T_ω. Equip $T_\omega \subset T$ with the usual trace σ-algebra $\mathcal{T} \cap T_\omega$, consisting precisely of all $\mathcal{T}$-measurable subsets of T_ω, and the usual trace measure $\mu(\cdot \cap T_\omega)$. Then by making the obvious substitutions in Definition 2.7, the set $\mathcal{S}_{\Sigma_\omega}$ can be equipped with its own feeble topology (and this will be done from now on). Let $U_\omega : \text{gph}\ \Sigma_\omega \times \mathcal{S}_{\Sigma_\omega} \to \mathbb{R}$ be the *utility function* of the game Γ_ω. Given the outcome $\omega \in \Omega$, player $t \in T_\omega$ faces payoff $U_\omega(t, s, f |_{T_\omega})$ if he/she takes action $s \in \Sigma_\omega(t)$ under the *a priori* action profile $f \in \mathcal{S}_\Sigma$.

Definition 2.1 For $\omega \in \Omega$ a *Nash equilibrium* for the game Γ_ω is an action profile $p \in \mathcal{S}_{\Sigma_\omega}$ such that

$$p(t) \in \text{argmax}_{s \in \Sigma(t)} U_\omega(t, s, p) \text{ for } \mu\text{-a.e. } t \text{ in } T_\omega.$$

Note that when T is finite and μ is the counting measure, the above "for μ-a.e. t in T_ω" amounts to "for all t in T_ω". Hence, the above definition extends Nash's classical notion.

Before stating our principal definition we recapitulate the essence of the game Γ. Recall that *all* players of T must make their action in ignorance of the random outcome ω to be realized under the probability P. In turn, this implies that the players are uncertain about the payoff and even about their participation in the game Γ_ω to be realized. After the realization of ω only the players in $T_\omega \subset T$ play the game Γ_ω and they are committed to the actions which they had chosen *a priori*. Similar situations occur in real life. For instance, consider the design and crew formation stage for a sailing-yacht competition, with Ω being the set of all possible weather conditions at the site of the match. More examples of this kind can be found in Borm *et al.* (1995) and Voorneveld (1999). Following these authors, we can now state the definition of MLNE. Shortly we will state a result which guarantees its sets A_f, $f \in \mathcal{S}_\Sigma$, to be measurable, so that this definition makes truly mathematical sense.

Definition 2.2 A *maximum likelihood Nash equilibrium* (MLNE) for Γ is an *a priori* action profile $f_* \in \mathcal{S}_\Sigma$ such that $L(f_*) = \sup_{f \in \mathcal{S}_\Sigma} L(f)$. Here the *equilibrium likelihood* $L(f)$ of $f \in \mathcal{S}_\Sigma$ is given by $L(f) := P(A_f)$, with

$$A_f := \{\omega \in \Omega : f |_{T_\omega} \text{ is a Nash equilibrium for } \Gamma_\omega\}$$

being the event that f constitutes a Nash equilibrium profile.

This MLNE-notion agrees with our introductory example, which had $\Omega = \mathbb{R}_+$, $T = \{1, 2\}$, $S = \{H, T\}$, $\Sigma(1) = \Sigma(2) = S$ and for instance $U_\omega(1, H, (f(1), f(2))) = 1 - \omega$ if $f(2) = H$ and $U_\omega(1, H, (f(1), f(2))) = -1 + 2\omega$ if $f(2) = T$, and so on. We now complete our list of measurability assumptions for Γ.

$$(\omega, t, s) \mapsto U_\omega(t, s, f |_{T_\omega}) \text{ is } E \cap (\mathcal{F} \otimes \mathcal{T} \otimes \mathcal{B}(S))\text{-measurable for every } f \in \mathcal{S}_\Sigma. \quad (2.4)$$

Define $V_\omega : T_\omega \times \mathcal{S}_{\Sigma_\omega} \to (-\infty, +\infty]$ by

$$V_\omega(t,p) := \sup_{s\in\Sigma(t)} U_\omega(t,s,p). \tag{2.5}$$

Then combining Definitions 2.1, 2.2 and (2.5) obviously gives

$$A_f = \{\omega \in \Omega : U_\omega(t, f(t), f\mid_{T_\omega}) = V_\omega(t, f\mid_{T_\omega}) \text{ for } \mu\text{-a.e. } t \text{ in } T_\omega\}. \tag{2.6}$$

We now establish that Definition 2.2 makes mathematical sense. Let $\hat{\mathcal{F}}$ be the P-completion of $\mathcal{F}$ and let $\hat{\mathcal{D}}$ be the completion of $\mathcal{D}$ with respect to the trace measure $(P\times\mu)(D\cap\cdot)$.

Theorem 2.3 *Under (2.1), (2.3) and (2.4) for every $f \in \mathcal{S}_\Sigma$ the function $(\omega,t)\mapsto V_\omega(t,f\mid_{T_\omega})$ is $\hat{\mathcal{D}}$-measurable and the set A_f in (2.6) is $\hat{\mathcal{F}}$-measurable.*

Lemma 2.4 *Let $\psi : D \to [0,+\infty]$ be a $\hat{\mathcal{D}}$-measurable function. Then there exists $F\in\mathcal{F}$ with $P(F)=1$ such that the following hold:*

(i) For every $\omega\in F$ the function $\psi(\omega,\cdot) : T_\omega \to [0,+\infty]$ is $\mathcal{T}$-measurable on T_ω.

(ii) The function $\omega\mapsto\int_{T_\omega}\psi(\omega,t)\mu(dt)$ from F into $[0,+\infty]$ is $\hat{\mathcal{F}}$-measurable.

Proof. *Step 1.* In this first step we assume that ψ is the characteristic function of a $\hat{\mathcal{D}}$-measurable subset $\hat{C}$ of D. By Proposition I.4.5 of Neveu (1965), $\hat{C}$ can be written as $\hat{C} = C\cup N$, where C belongs to $\mathcal{D}$ and $N\subset D$ is $P\times\mu$-negligible. So then N is contained in some $B\in\mathcal{D}$ with $(P\times\mu)(B)=0$. This means that for P-a.e. ω in Ω the section B_ω of B at ω has μ-measure zero by Corollary 1 on p. 76 of Neveu (1965). In other words, for all ω in a set $F\in\mathcal{F}$ of P-measure 1 one has $\mu(B_\omega)=0$. For every $\omega\in F$ the following holds: the inclusions $C_\omega\subset\hat{C}_\omega\subset C_\omega\cup B_\omega$, with $\mu(\hat{C}_\omega)=\mu(C_\omega)$, imply that $\hat{C}_\omega$ belongs to the μ-completion of $\mathcal{T}$, which is $\mathcal{T}$ itself (recall that $(T,\mathcal{T},\mu)$ is complete). Since $\hat{C}_\omega\subset T_\omega$ is obvious, we have that $\hat{C}_\omega$ belongs to $T_\omega\cap\mathcal{T}$. Also, the previous results give $\mu(\hat{C}_\omega)=\mu(C_\omega)$ for all $\omega\in F$, and $\omega\mapsto\mu(C_\omega)$ is $\mathcal{F}$-measurable on Ω by Corollary 2 on p. 76 of Neveu (1965). It remains to observe that $\psi(\omega,\cdot)=1_{\hat{C}_\omega}$ for all $\omega\in\Omega$ and that $\int_{T_\omega}\psi(\omega,t)\mu(dt)=\mu(C_\omega)$ for all $\omega\in F$.

Step 2. Next, suppose that ψ is of the form $\psi=\sum_{i=1}^m c_i 1_{\hat{C}^i}$, with $c_1,\dots,c_m$ nonnegative constants and $\hat{C}^1,\dots,\hat{C}^m$ in $\hat{\mathcal{D}}$. By the previous step it is easy to see that $\psi(\omega,\cdot)=\sum_i c_i 1_{\hat{C}^i_\omega}$ is $\mathcal{T}$-measurable for all ω in a set $F\in\mathcal{F}$ of full P-measure. So on that same set the function $\omega\mapsto\int_{T_\omega}\psi(\omega,t)\mu(dt)$ coincides with a finite sum of $\mathcal{F}$-measurable functions, that is to say, with a $\mathcal{F}$-measurable function.

Step 3. Finally, if ψ is nonnegative and $\hat{\mathcal{D}}$-measurable, then ψ is pointwise a monotone limit of step functions of the type considered in step 2. So by step 2 and an obvious argument involving pointwise limits, $\psi(\omega,\cdot)$ is $\mathcal{T}$-measurable for P-a.e. ω. This proves (i), and (ii) follows by an application of the monotone convergence theorem QED

Proof of Theorem 2.3. Fix $f \in \mathcal{S}_\Sigma$ arbitrarily. By (2.5) we have For any $\alpha \in \mathbb{R}$

$$V_\omega(t, f\mid_{T_\omega}) > \alpha \text{ if and only if there exists } s \in \Sigma(t) \text{ with } U_\omega(t, s, f\mid_{T_\omega}) > \alpha.$$

Hence, the set C_α of all $(\omega, t) \in D$ with $V_\omega(t, f\mid_{T_\omega}) > \alpha$ is the projection onto D of the set of all $(\omega, t, s) \in E$ for which $U_\omega(t, s, f\mid_{T_\omega}) > \alpha$. By (2.4) the latter set belongs to the trace σ-algebra $E \cap (\mathcal{F} \otimes \mathcal{T} \otimes \mathcal{B}(S))$, whence to $(D \cap (\mathcal{F} \otimes \mathcal{T})) \otimes \mathcal{B}(S)$, i.e., to $\mathcal{D} \otimes \mathcal{B}(S)$. Since S is a Suslin space, the measurable projection Theorem III.23 in Castaing and Valdier (1977) guarantees that C_α is a universally measurable subset of $(D, \mathcal{D})$. This means in particular that C_α is $\hat{\mathcal{D}}$-measurable. So $(\omega, t) \mapsto V_\omega(t, f\mid_{T_\omega})$ has been shown to be $\hat{\mathcal{D}}$-measurable. Define now the $\hat{\mathcal{D}}$-measurable function $\psi : D \to [0, +\infty]$ by

$$\psi(\omega, t) := V_\omega(t, f\mid_{T_\omega}) - U_\omega(t, f(t), f\mid_{T_\omega}).$$

By Lemma 2.4 there exists a set F in $\mathcal{F}$, $P(F) = 1$, such that for every $\omega \in F$ the function $\psi(\omega, \cdot) : T_\omega \to [0, +\infty]$ is $T_\omega \cap \mathcal{T}$-measurable and the mapping $\omega \mapsto \int_{T_\omega} \psi(\omega, t)\mu(dt)$ is measurable with respect to $F \cap \hat{\mathcal{F}}$. Since

$$A_f \cap F = \{\omega \in F : \int_{T_\omega} \psi(\omega, t)\mu(dt) = 0\},$$

it follows that $A_f \cap F$ is $\hat{\mathcal{F}}$-measurable. But then so is A_f, since $A_f \cap (\Omega \backslash F)$ is P-negligible. QED

Next, we move to the issue of MLNE existence. The following example shows that even for one-player games such existence does not hold without further assumptions.

Example 2.5 Let $\Omega := \mathbb{N}$, $S := \mathbb{N}$, $\Sigma \equiv S$ and let T be the singleton $\{1\}$. Then $\mathcal{S}_\Sigma = \mathbb{N}$. Let P be given by $P(\{n\}) := 2^{-n}$, $n \in \mathbb{N}$. Also, let $U_\omega(1, s, f) := -1$ if $s = \omega$ and $U_\omega(1, s, f) := 0$ otherwise (this means that the player wishes to guess ω *incorrectly*). Then $L(n) = P(\mathbb{N}\backslash\{n\}) = 1 - 2^{-n}$. Evidently, this game does not have an MLNE at all.

We list our topological and geometrical conditions under which existence of an MLNE holds (cf. Balder (2000c)). For the multifunction $\Sigma : T \to 2^S$ we make the following assumptions:

$$\Sigma(t) \text{ is nonempty and compact for every } t \in T^{pa}, \tag{2.7}$$

$$\Sigma(t) \text{ is nonempty, convex and compact for every } t \in T^{na}. \tag{2.8}$$

Under these assumptions it follows from our earlier remarks about the weak metric ρ that the ρ-topology coincides with the original S-topology on the compact sets $\Sigma(t)$, $t \in T$.

Remark 2.6 Given the above assumptions, we notice that under the following additional condition

$$U_\omega(t,\cdot,p) \text{ is lower semicontinuous on } \Sigma(t) \text{ for every } (t,\omega,p) \in D \times \mathcal{S}_{\Sigma_\omega},$$

which is amply fulfilled in Borm *et al.* (1995) and Voorneveld (1999), the following identity holds:

$$V_\omega(t, f\mid_{T_\omega}) = \sup_i U_\omega(t, s_i(t), f\mid_{T_\omega}).$$

Here the countable collection (s_i) in $\mathcal{S}_\Sigma$ is such that $\{s_i(t) : i \in \mathbb{N}\}$ is dense in $\Sigma(t)$ for every $t \in T$; i.e., (s_i) is a Castaing representation. This follows from applying Theorem III.7 in Castaing and Valadier (1977), which is made possible by virtue of (2.3) and the above observation about the weak metric ρ. Hence, under such additional lower semicontinuity the measurability result in Theorem 2.3 becomes an immediate consequence of (2.4) that neither requires the measurable projection theorem nor the Suslin property of S.

From now on, the set of profiles $\mathcal{S}_\Sigma$ will be equipped with the feeble topology, which is defined as follows. Let $\mathcal{G}_{LC,\Sigma}$ be the vector space of all $\mathcal{T} \times \mathcal{B}(S)$-measurable $g : T \times S \to \mathbb{R}$ such that $g(t,\cdot)$ is linear and continuous on S for every $t \in T$, and such that $\sup_{s\in\Sigma(t)} |g(t,s)| \leq \phi_g(t)$ for all $t \in T$ for some μ-integrable function $\phi_g \in \mathcal{L}^1_{\mathbb{R}}(T,\mathcal{T},\mu)$. Observe that for $g \in \mathcal{G}_{LC,\Sigma}$ the integral expression

$$J_g(f) := \int_T g(t, f(t))\mu(dt)$$

is well-defined for every $f \in \mathcal{S}_\Sigma$. The following definition was introduced by Balder (1999b). See Meyer (1973) for a similar notion of "median limits", formulated in quite different terms and based on the continuum hypothesis.

Definition 2.7 The *feeble topology* on $\mathcal{S}_\Sigma$ is the coarsest topology for which the integral functionals $J_g : \mathcal{S}_\Sigma \to \mathbb{R}$ are continuous for all $g \in \mathcal{G}_{LC,\Sigma}$.

The versatile nature of this topology is illustrated by the following. Example 2.1 in Balder (1999b), which has $(T,\mathcal{T},\mu)$ in addition separable, shows that if S is a separable Banach space and Σ is μ-integrably bounded, then the feeble topology on $\mathcal{S}_\Sigma$ coincides with the usual weak $\mathcal{L}^1$-topology. Also, Example 2.2 in Balder (1999b) shows that if S is the dual of a separable Banach space and if Σ is actually uniformly bounded, then the feeble topology on $\mathcal{S}_\Sigma$ coincides with the usual weak star $\mathcal{L}^\infty$-topology. Our remaining assumptions for the payoff structure are as follows:

$$U_\omega(t,\cdot,\cdot) \text{ is upper semicontinuous on } \Sigma(t) \times \mathcal{S}_{\Sigma_\omega} \text{ for every } (\omega,t) \in D, \tag{2.9}$$

$$\text{argmax}_{s\in\Sigma(t)} U_\omega(t,s,p) \text{ is convex for every } (\omega,t,p) \in D \times \mathcal{S}_{\Sigma_\omega} \text{ with } t \in T^{na}, \tag{2.10}$$

$$V_\omega(t,\cdot) \text{ is lower semicontinuous on } \mathcal{S}_{\Sigma_\omega} \text{ for every } t \in T_\omega. \tag{2.11}$$

Observe that (2.10) holds in particular if $U_\omega(t,\cdot,f)$ is quasiconcave for every $(\omega,t) \in D$, $t \in T^{na}$, and for every $f \in \mathcal{S}_{\Sigma_\omega}$. Observe also that (2.9) and (2.11) hold in particular if $U_\omega(t,\cdot,\cdot)$ is continuous on $\Sigma(t) \times \mathcal{S}_{\Sigma_\omega}$ for every $(\omega,t) \in D$.

Theorem 2.8 *Under (2.1), (2.3), (2.4) and (2.7)–(2.11) there exists an* a priori *action profile that is an MLNE for the random continuum game* Γ.

The proof of this theorem will now proceed by means of Young measure theory; cf. Balder (2000a, 2000b). It is composed of several lemmas. Our approach is reminiscent of the classical direct method in the calculus of variations, but it uses generalized limits of the minimizing sequence in the form of transition probabilities, i.e., Young measures. First, observe that $\mathcal{S}_\Sigma$ contains at least one element $\tilde{f}$ by the von Neumann-Aumann measurable selection Theorem III.22 in Castaing and Valadier (1977), in view of (2.3), (2.7) and (2.8). Because of this nonemptiness of $\mathcal{S}_\Sigma$, there certainly exists a sequence (f_n) in $\mathcal{S}_\Sigma$ such that $\lim_n L(f_n) = \sup_{f \in \mathcal{S}_\Sigma} L(f)$. Let $\mathcal{R}_S$ stand for the set of all transition probabilities δ with respect to $(T,\mathcal{T})$ and $(S,\mathcal{B}(S))$. Also, let $\epsilon_f \in \mathcal{R}_S$ be the canonical transition probability associated with the function $f \in \mathcal{S}_\Sigma$.

Lemma 2.9 *There exist a subsequence (f_m) of (f_n) and a $\delta_* \in \mathcal{R}_S$ for which*

$$(\epsilon_{f_m}) \textit{ converges narrowly in } \mathcal{R}_S \textit{ to } \delta_*,$$

$$\delta_*(t)(\Lambda_t) = 1 \textit{ for } \mu\textit{-a.e. } t \textit{ in } T$$

and

$$\lambda_t := \lim_m f_m(t) \textit{ exists for } \mu\textit{-a.e. } t \textit{ in } T^{pa}.$$

Here $\Lambda_t \subset \Sigma(t)$ denotes the set of all limit points of $(f_m(t))$.

For the definition of narrow convergence (alias Young measure convergence) we refer for instance to Balder (1988) or to Definition 4.1 in Balder (2000a).

Proof. On every atom A_j constituting T^{pa} every function f_n is constant μ-a.e. Also, Σ is constant μ-a.e. to a compact subset Σ_j of S on such an atom A_j (use Castaing representation as in Theorem III.7 of Castaing and Valadier (1977)). So by a diagonal extraction argument we obtain a preliminary subsequence (f_k) of (f_n) for which the third property holds (recall that the weak metric ρ can be used on the sets Σ_j). Thanks to (2.7), the first result now follows directly by an application to (ϵ_{f_k}) of Prohorov's theorem for Young measures (cf. Theorem 4.10 of Balder (2000a) or Proposition 3.1 of Balder (1999b)). Notice that the third property continues to hold for the further subsequence (f_m) of (f_k) that is obtained in this way. Also, the second result follows from the first one by the support theorem for narrow convergence (see Theorem 4.12 of Balder (2000a)). This second result holds as stated thanks to the fact that the sets $\Sigma(t)$, $t \in T$, can be equipped with the equivalent weak metric ρ; this was mentioned immediately following (2.7)–(2.8). QED

Obviously, the second and third property of Lemma 2.9 imply that Λ_t is the singleton $\{\lambda_t\}$ for μ-a.e. t in the purely atomic part T^{pa}. Also, the second property in Lemma 2.9 implies that $\delta_*(t)(\Sigma(t)) = 1$ for μ-a.e. t in T. So it follows by (2.7)–(2.8) that the barycenter

$$f_\infty(t) := \text{bar } \delta_*(t) := \int_S s\delta_*(t)(ds) \tag{2.12}$$

is well-defined for μ-a.e. t, i.e., for all t outside a certain μ-null set N (apply Proposition 26.3 of Choquet (1969)). Setting $f_\infty(t) := \tilde{f}(t)$ on $T \backslash N$ thus creates an element f_∞ of $\mathcal{S}_\Sigma$ (recall that $\tilde{f}$ was obtained above to show that $\mathcal{S}_\Sigma$ is nonempty). Observe that we must have

$$f_\infty(t) = \lambda_t := \lim_m f_m(t) \text{ for } \mu\text{-a.e. } t \text{ in } T^{pa} \tag{2.13}$$

by the third part of Lemma 2.9.

Lemma 2.10 *The sequence (f_m) converges to f_∞ in $\mathcal{S}_\Sigma$ in the feeble topology.*

Proof. A direct application of Proposition 3.2 of Balder (1999b). QED

Lemma 2.11 *For every $\omega \in \Omega$ the mapping $f \mapsto f\,|_{T_\omega}$ from $\mathcal{S}_\Sigma$ to $\mathcal{S}_{\Sigma_\omega}$ is continuous for the respective feeble topologies.*

Proof. An elementary consequence of Definition 2.7 and (2.1); simply observe that for every $\omega \in \Omega$ and $g \in \mathcal{G}_{LC,\Sigma_\omega}$ the function $\tilde{g} : T \times S \to \mathbb{R}$ defined by

$$\tilde{g}(t,s) := \begin{cases} g(t,s) & \text{if } t \in T_\omega, \\ 0 & \text{otherwise} \end{cases}$$

belongs to $\mathcal{G}_{LC,\Sigma}$. QED

Lemma 2.12 *For every $\omega \in \Omega$ the function $\ell_\omega : \text{gph } \Sigma_\omega \times (\mathbb{N} \cup \{\infty\}) \to [0, +\infty]$, defined by*

$$\ell_\omega(t,s,m) := V_\omega(t, f_m\,|_{T_\omega}) - U_\omega(t,s,f_m\,|_{T_\omega}),$$

has the following properties:

$$\ell_\omega(\cdot,\cdot,m) \textit{ is } \text{gph } \Sigma \cap (\mathcal{T} \otimes \mathcal{B}(S))\textit{-measurable for every } m \in \mathbb{N},$$

$$\ell_\omega(t,\cdot,\cdot) \textit{ is lower semicontinuous on } \Sigma(t) \times (\mathbb{N} \cup \{\infty\}) \textit{ for every } t \in T_\omega.$$

Here the usual Alexandrov topology is used on $\mathbb{N} \cup \{\infty\}$.

Proof. Similar to the proof of Theorem 2.3, it follows by the measurable projection Theorem III.23 of Castaing and Valadier (1977) that $t \mapsto V_\omega(t, f_m\,|_{T_\omega})$ is $\mathcal{T}$-measurable (observe that $\mathcal{T}$ is μ-complete). In view of (2.4), this proves the first property. The second property follows from (2.9)–(2.11) and Lemma 2.10 (note that $(f_m\,|_{T_\omega})$ feebly converges to $f_\infty\,|_{T_\omega}$ by Lemma 2.11). QED

Lemma 2.13 $\limsup_{m\to\infty} L(f_m) \leq L(f_\infty)$.

PROOF. Recall from Definition 2.2 that $L(f) = P(A_f)$ for every $f \in \mathcal{S}_\Sigma$. Since $\limsup_m P(A_{f_m}) \leq P(\limsup_m A_{f_m})$, the stated result follows if we can prove

$$\limsup_m A_{f_m} \subset A_{f_\infty}. \tag{2.14}$$

To this end, fix $\omega \in \limsup_m A_{f_m} := \cap_{m' \in \mathbb{N}} \cup_{m \geq m'} A_{f_m}$ arbitrarily. Then $\omega \in A_{f_m}$ for infinitely many m, so $\liminf_m \int_{T_\omega} \ell_\omega(t, f_m(t), m)\mu(dt) = 0$. By the lower closure Theorem 4.13 in Balder (2000a) (take its D equal to $\mathbb{N} \cup \{\infty\}$, equipped with the Alexandrov topology, and set its functions d_m identically equal to m), it follows from Lemmas 2.9 and 2.12 that

$$\liminf_m \int_{T_\omega} \ell_\omega(t, f_m(t), m)\mu(dt) \geq \int_{T_\omega} [\int_{\Sigma(t)} \ell_\omega(t, s, \infty)\delta_*(t)(ds)]\mu(dt),$$

and this gives

$$\int_S \ell_\omega(t, s, \infty)\delta_*(t)(ds)\mu(dt) = 0 \text{ for } \mu\text{-a.e. } t \text{ in } T_\omega. \tag{2.15}$$

For any non-exceptional t we now claim that

$$f_\infty(t) \in M_\infty(t) := \text{argmax}_{s \in \Sigma(t)} U_\omega(t, s, f_\infty). \tag{2.16}$$

If $t \in T^{pa}$, then Lemma 2.9 and (2.13) imply that $\delta_*(t)$ is concentrated at the point $f_\infty(t)$. Thus, (2.15) amounts to $\ell_\omega(t, f_\infty(t), \infty) = 0$, i.e., to (2.16). On the other hand, if $t \in T^{na}$, then (2.15) implies for $\delta_*(t)$-almost all $s \in \Sigma(t)$ that $\ell_\omega(t, s, \infty) = 0$, whence $s \in M_\infty(t)$. This means that $f_\infty(t)$, the barycenter of $\delta_*(t)$, lies in the closed convex hull of $M_\infty(t)$ (apply the Hahn-Banach theorem). But assumptions (2.9) and (2.10) cause $M_\infty(t)$ to be closed and convex, so (2.16) holds also in this case. This proves the claim. We thus conclude that (2.16) holds for μ-almost all $t \in T_\omega$. This implies $\omega \in A_{f_\infty}$ and therefore (2.14) has been proven. QED

PROOF OF THEOREM 2.8. Evidently, our choice of the original sequence (f_n) implies that the subsequence (f_m) satisfies $\sup_{f \in \mathcal{S}_\Sigma} L(f) = \lim_m L(f_m)$. Hence, it follows from Lemma 2.13 that $f_\infty \in \mathcal{S}_\Sigma$, defined in (2.12), is a maximum likelihood equilibrium profile. QED

Our next remark shows that ω-dependence of the feasible action sets can easily be absorbed by the technically more simple model of this paper.

Remark 2.14 Consider the following extension $\tilde{\Gamma}$ of the game Γ. We are given multifunctions $\tilde{\Sigma}_\omega$ from T_ω into 2^S, $\omega \in \Omega$, with $\tilde{\Sigma}_\omega(t) \subset \Sigma(t)$ for all $t \in T_\omega$. In $\tilde{\Gamma}$, if the outcome $\omega \in \Omega$ is realized μ-almost every player $t \in T_\omega$ is forced to take his/her action in the subset $\tilde{\Sigma}_\omega(t)$ of $\Sigma(t)$. Thus, in Definition 2.2 the *a priori* action profile $f \in \mathcal{S}_{\tilde{\Sigma}}$ is only deemed to be a Nash equilibrium for $\tilde{\Gamma}_\omega$, the associated realization of $\tilde{\Gamma}$, if

$$f(t) \in \text{argmax}_{s \in \tilde{\Sigma}_\omega(t)} U_\omega(t, s, f\,|_{T_\omega}) \text{ for } \mu\text{-a.e. } t \text{ in } T_\omega. \tag{2.17}$$

It is easy to absorb this additional feature in the model of this paper by introducing the following new payoffs:

$$\tilde{U}_\omega(t,s,p) := \begin{cases} \arctan U_\omega(t,s,p) & \text{if } s \in \tilde{\Sigma}_\omega(t) \\ -2 & \text{if } s \in \Sigma(t)\backslash\tilde{\Sigma}_\omega(t) \end{cases}$$

Indeed, this follows immediately from the obvious identity

$$\mathrm{argmax}_{s\in\Sigma(t)}\tilde{U}_\omega(t,s,f\mid_{T_\omega}) = \mathrm{argmax}_{s\in\tilde{\Sigma}_\omega(t)}U_\omega(t,s,f\mid_{T_\omega}), \tag{2.18}$$

which holds for all $(\omega,t) \in D$ and all $f \in \mathcal{S}_\Sigma$. Under the following analogue of (2.1), (2.3) and (2.4)

$$\tilde{E} := \{(\omega,t,s) \in \Omega\times T\times S : s \in \tilde{\Sigma}_\omega(t)\} \text{ is } \mathcal{F}\otimes\mathcal{T}\otimes\mathcal{B}(S)\text{-measurable},$$

$$(\omega,t,s) \mapsto U_\omega(t,s,f\mid_{T_\omega}) \text{ is } \tilde{E}\cap(\mathcal{F}\otimes\mathcal{T}\otimes\mathcal{B}(S))\text{-measurable for every } f\in\mathcal{S}_\Sigma,$$

and under the following analogues of (2.7)–(2.8)

$$\tilde{\Sigma}_\omega(t) \text{ is nonempty and compact for every } \omega\in\Omega \text{ and } t\in T^{pa}\cap T_\omega,$$

$$\tilde{\Sigma}_\omega(t) \text{ is nonempty, convex and compact for every } \omega\in\Omega \text{ and } t\in T^{pa}\cap T_\omega,$$

together with the obvious analogues of (2.9)–(2.11), obtained by replacing $\Sigma(t)$ by $\tilde{\Sigma}_\omega(t)$, it is easy to see that that Theorem 2.8 continues to hold for the game $\tilde{\Gamma}$. This results in the existence of an MLNE for $\tilde{\Gamma}$ (note that it turns out to be an MLNE in the sense of Definition 2.2 and (2.17), because of (2.18)).

Remark 2.15 Let $\Sigma_j \subset S$ be as in the proof of Lemma 2.9. It is easy to see that the set $\mathcal{S}' := \{f\mid_{T^{pa}}: f\in\mathcal{S}_\Sigma\}$ can be identified with the Cartesian product $\Pi_j\Sigma_j$ and that the restriction of the feeble topology to $\mathcal{S}'$ coincides with the usual product topology on $\Pi_j\Sigma_j$. Also, the proof of the third property in Lemma 2.9 actually requires no linear structure for S. These facts imply that for the players in every atom A_j of T^{pa} we could have worked with an action space S_j that is metrizable and Suslin (or simply metrizable and compact by setting $\Sigma_j = S_j$), but not necessarily linear. On the nonatomic part T^{na} we would still keep to the original Suslin locally convex vector space S.

Let us compare Theorem 2.8 with the results of Borm *et al.* (1995) and Voorneveld (1999). Both these papers have $T^{na} = \emptyset$, $T^{pa} = T$. Their T is a finite set, so one can use the power set 2^T as $\mathcal{T}$ and the counting measure as μ. As already observed by Voorneveld (1999), the setup of Borm *et al.* (1995) is incomplete: the proof of their Lemma 1 contains errors and, in order to go through, additional upper semicontinuity of the payoff functions seems indispensible. If indeed such upper semicontinuity is added to their conditions, then our Theorem 2.8 considerably generalizes their main existence result. To see this, we observe that the payoffs in that paper are supposed to be lower semicontinuous in the actions (hence Remark 2.6 applies). Together with the upper semicontinuity required to repair their result, above, this causes their payoffs to be continuous in the actions. Hence, product measurability of the payoffs in Borm *et al.* (1995) follows from Lemma III.14 of Castain and Valadier (1977). In fact, Remark 2.15 then shows that it is enough to work with compact metric spaces of actions S_j for the players in each atom A_j. Going now even one step further, because of the dispensability of the Suslin property (cf. Remark 2.6): by concentrating on the restrictions to T^{pa} of the minimizing sequence (f_n), as used in the proof of Theorem 2.8, we can simply require those spaces S_j to be sequentially compact to retain the third property of Lemma 2.9.[1] Voorneveld (1999) works with preference relations rather than payoffs, but the conditions on these (see his pp. 219-220) are such that they allow the utility-type representation as used in this paper (apply Theorem 3.1 of Jehle and Reny (1998)). Remark 2.14 applies to his model; see in particular his p. 223. Just as in Borm *et al.* (1995), which is simpler by not allowing for random participation, only a finite set T of players is considered by Voorneveld (1999) (as mentioned before, this is a purely atomic measure space for the power set σ-algebra and the counting measure). Also, the space S of all possible actions is Euclidean and the multifunction Σ is constant (see his pp. 224, 227)

Notice also that the proof of Theorem 2.8 given here is quite different from the one given in Balder (2000c), which offers a viable alternative that substitutes new results about the feeble topology for the present paper's reliance on Young measure theory. That paper also uses another approach as a whole, which is based on the abstract scheme given in Proposition 2.1 of Balder (2000c).

It is also interesting to compare Theorem 2.8 to the usual existence results for Nash equilibria in continuum game theory. The differences are as follows. Obviously, the existence of *a priori* action profiles $f_* \in \mathcal{S}_\Sigma$ such that $f_* \mid_{T_\omega}$ is a Nash equilibrium of Γ_ω for P-almost every ω (i.e., $L(f_*) = 1$), or even with $L(f_*) > 0$, is out of the question in even the simplest cases. For instance, if we allow mixing in the example given in the introduction, then we can work with $\Sigma(1) = \Sigma(2) = [0,1]$ in an obvious way (i.e., for $t = 1, 2$ we let $f(t) \in \Sigma(t)$ stand for the probability measure that assigns probability $f(t)$ to "heads" and $1 - f(t)$ to "tails"). This gives

$$U_\omega(t, s, (f_1, f_2)) := \begin{cases} 2s(2f_2 - 2\omega f_2 - 1 + \omega) & \text{if } t = 1, \\ s(-4f_1 + 4\omega f_1 + 2 - 3\omega) & \text{if } t = 2. \end{cases}$$

For instance, for $\omega \in [0, \frac{2}{3}) \cup (2, +\infty)$ it follows elementarily that the ω-dependent mixed action $(s_1^*, s_2^*) = (\frac{3\omega - 2}{4\omega - 4}, \frac{1}{2})$ constitutes the unique mixed Nash equilibrium for the game Γ_ω. Thus, for any nonatomic probability measure P that is entirely concentrated on $[0, \frac{2}{3}) \cup (2, +\infty)$, we have $L(f_1, f_2) = 0$ for all a priori action profiles (f_1, f_2). To make the comparison more fair, one can restrict onseself for instance to the situation where ω is *fixed*. It then turns out that to guarantee existence of a Nash equilibrium for the game Γ_ω, with $\omega \in \Omega$ fixed, the following conditions must hold in addition to (2.1)–(2.11) (cf. Balder (1995, 1999b, 2001)): For every $t \in T^{pa} \cap T_\omega$ two additional convexity conditions must hold: (i) $\Sigma(t)$

[1] In the proof of Theorem 1 of Borm *et al.* (1995) sequential compactness, rather than compactness itself, has to be used.

is convex, (ii) the set $\text{argmax}_{s\in\Sigma(t)}U_\omega(t,s,p)$ is convex for every $p \in \mathcal{S}_{\Sigma_\omega}$. This need for additional convexity forms an essential difference with Theorem 2.8. For instance, by fixing $\omega = 0$ in the original example in the introduction we get the classical "matching pennies" game, for which no pure Nash equilibrium exists. But for P equal to the Dirac probability measure at 0 we already noticed that four MLNE's in pure action profiles exist vacuously.

Another difference with the usual existence literature is that Theorem 2.8 requires convexity of the sets $\Sigma(t)$ and $\text{argmax}_{s\in\Sigma(t)}U_\omega(t,s,f)$ for t in the nonatomic part T^{na} (see (2.8) and (2.10)). However, consider that not only a very general utility function is used here, but also that, if one were to introduce more special versions of the utility function, the usual purification by nonatomicity runs into the obstacle of ω-dependence of T_ω. However, in case participation is not random, one could certainly use the usual purification by nonatomicity; cf. Balder (1995, 2001). The following example illustrates this.

Example 2.16 Let Ω be the singleton $\{0\}$ and let $T := [0,1]$ be equipped with the Lebesgue σ-algebra and the Lebesgue measure λ. Let $S := [-1,1]$, $\Sigma \equiv S$ and set $\mu := \lambda$. Also, let $U_0(t,s,f) := -s\int_0^1 f\,d\lambda$. It is easy to see that $f_* \in \mathcal{S}_\Sigma$ constitutes a Nash equilibrium (and an MLNE as well) if and only if $\int_0^1 f_*\,d\mu = 0$ for $\mu = \lambda$. So such f_* exist in abundance. In contrast, suppose next that in the same setting one works with a measure $\mu := \lambda + 2\epsilon_1$ (here ϵ_1 is the point probability at the point 1) and that player 1's feasible action set is restricted to be the nonconvex set $\Sigma(1) := \{-1,1\}$ (all other players continue to have $[-1,1]$ as their feasible action space). Just as above, $f_* \in \mathcal{S}_\Sigma$ is NE if and only if $\int_0^1 f_* d\mu = 0$. It is easy to check that no such f_* exists. So there does not exist a Nash equilibrium for this second version of the game. But of course, in both cases this game has an MLNE. Observe finally that when the nonatomic player 1 is allowed to take *mixed* actions, Nash equilibria exist once again.

Acknowledgment. A part of this research was done during a stay at the Department of Economics of the University of Alabama in Tuscaloosa (USA). I thank the department for its hospitality and support.

References.

Aumann R (1964) Markets with a continuum of traders. Econometrica **32**:39-50.

Balder EJ (1988) Generalized equilibrium results for games with incomplete information. Math Oper Res **13**:265-276.

Balder EJ (1995) A unifying approach to existence of Nash equilibria. Intern J Game Theory **24**:79-94.

Balder EJ (1999a) Young measure techniques for existence of Cournot-Nash-Walras equilibria. In: Wooders, M (ed.) Topics on Mathematical Economics and Game Theory: Essays in honor of R.J. Aumann. Fields Institute Communications No. 23, American Mathematical Society, Providence, RI, pp. 31–39.

Balder EJ (1999b) On the existence of Cournot-Nash equilibria in continuum games. J Math Econ **32**, 207-223.

Balder EJ (2000a) New fundamentals of Young measure convergence," in: A. Ioffe, S. Reich and I. Shafrir (eds.) Calculus of Variations and Related Topics, Research Notes in Mathematics no. 411. Chapman and Hall, Boca Raton, pp. 24-48.
Balder EJ (2000b) Lectures on Young measure theory and its applications in economics. Rend Ist Mat Univ Trieste **31**, Suppl. 1:1-69.
Balder EJ (2000c) On the existence of maximum likelihood Nash equilibria. Submitted to Annals Oper Res (special issue), June 2000.
Balder EJ (2001) A unifying pair of Cournot-Nash equilibrium results. J Econ Theory, to appear.
Borm PEM, Cao R and García-Jurado I (1995) Maximum likelihood equilibria of random games. Optimization **35**:77-84.
Castaing C and Valadier M (1977). Convex Analysis and Measurable Multifunctions. Lecture Notes in Math. No. 580, Springer-Verlag, Berlin.
Choquet G (1969) Lectures on Analysis. Benjamin, Reading, MS.
Dellacherie C and Meyer P-A (1975) Probabilités et Potentiel. Hermann, Paris. (English translation: North-Holland, Amsterdam, 1978).
Dudley RM (1989) Real Analysis and Probability. Wadsworth, Pacific Grove, CA.
Jehle GA and Reny P (1998) Advanced Microeconomic Theory. Addison-Wesley, New York.
Meyer P-A (1973) Limites médiales, d'après Mokobodzki. in: Séminaire de Probabilités VII, Lecture Notes in Math. No. 321, Springer-Verlag, Berlin, pp. 198-204.
Neveu J (1965) Mathematical Foundations of the Calculus of Probability. Holden-Day, San Fransisco.
Schmeidler D (1973) Equilibrium points of non-atomic games. J Statist Phys **7**:295-300.
Schwartz L (1975) Radon Measures. Oxford University Press, Oxford.
Voorneveld M (1999) Potential Games and Interactive Decisions with Multiple Criteria. Dissertation Series No. 61. Center for Economic Research, Tilburg University, Tilburg.