

Perfect equilibrium and lexicographic beliefs

Srihari Govindan and Tilman Klumpp

Department of Economics, University of Western Ontario, N6A 5C2 London, Ontario,
CANADA (email: sgovinda@uwo.ca)

Department of Economics, University of Western Ontario (email: tklumpp@uwo.ca)

Received February 2002/Final version June 2002

Abstract. We extend the results of Blume, Brandenberger, and Dekel (1991b) to obtain a finite characterization of perfect equilibria in terms of lexicographic probability systems (LPSs). The LPSs we consider are defined over individual strategy sets and thus capture the property of independence among players' actions. Our definition of a product LPS over joint actions of the players is shown to be canonical, in the sense that any independent LPS on joint actions is essentially equivalent to a product LPS according to our definition.

Key words: Perfect equilibrium; lexicographic beliefs; independence; finite characterization.

1. Introduction

The central idea in the literature on refinements of Nash equilibria is Selten's (1975) concept of a perfect equilibrium. A mixed strategy profile is perfect if it is a best reply against each profile in a sequence of completely mixed strategy profiles converging to it. For normal-form games, perfection strengthens the admissibility criterion used in decision theory and is in fact equivalent to it for two-player games. Admissibility – the requirement that players not use weakly dominated strategies – is justified on the grounds that no pure strategy of a player is considered a null event (in the sense of Savage) by his opponents in making their decisions. Formally, a strategy is admissible if it is a best reply against a completely mixed (possibly correlated) strategy profile. Perfection

We would like to thank Adam Brandenburger, Eddie Dekel, Arthur Robson, Peter Streufert, and Robert Wilson for their helpful comments. Govindan acknowledges research support from the Social Sciences and Humanities Research Council of Canada.

imposes the additional restriction that there is a sequence of such strategy profiles, each uncorrelated across players, that converges to the equilibrium.

The Selten programme remained incomplete to the extent that a finite characterization of perfection was not obtained. Because the definition involves universal and existential quantifiers, a direct check for whether a strategy profile is a perfect equilibrium requires, in principle, an infinite number of computations. Yet, the set of perfect equilibria is semi-algebraic (cf. Blume and Zame, 1994) and, therefore, it admits a finite characterization. In this paper, we obtain a finite characterization in terms of lexicographic probability systems (LPSs).

Blume, Brandenberger, and Dekel (1991b; henceforth, BBD) show that any convergent sequence of strategy profiles induces an LPS. They use this relation between sequences and LPSs to provide a characterization of perfect equilibria in terms of the latter. However, their characterization is not finite and, in particular, does not assure a product form that represents independence across players. In this paper, we continue the work of BBD to obtain an exact finite characterization using product LPSs.

The problematic feature of the BBD characterization of perfect equilibria is that it uses LPSs defined directly on the space of profiles of pure strategies for all players simultaneously. A possible interpretation of their result is that an LPS over strategy profiles represents an outside observer's assessment, from which beliefs of players concerning the strategies of their opponents are then derived.¹ The marginal LPS on a single player's strategy set might thus be viewed as the common belief of all other players about this player's strategy, if indeed it is independent of the others'. However, there is no way to compose the marginal LPSs to recover the joint LPS if independence is violated. Hence, it is unclear to what extent the various marginals can be considered independent. Further, there is no finite procedure to check whether an LPS over profiles of joint actions is induced by a sequence of mixed strategies.

From a conceptual viewpoint, it is desirable to obtain an equivalent formulation of perfection using LPSs over individual strategy sets that are composed to obtain beliefs over joint actions. Indeed, if players act independently, as is commonly assumed in noncooperative game theory, then alternative theories of how play would proceed, as described in LPSs, should be constructed from a hierarchy of alternative theories about how each individual player would play if he were not to play his equilibrium strategy. Thus, the primitives should be LPSs over each player's pure strategies, not LPSs over profiles of pure strategies.

From a practical viewpoint too, it is desirable to describe perfect equilibria using individual LPSs: in most applications, selections from the perfect equilibria of a model are made based on the plausibility of the beliefs supporting them, which in turn depend on the reasonableness of the alternate strategy choices that are posited for each player individually. An important example is use of the principle of forward induction as an equilibrium selection criterion.

Our aim, then, is to provide a definition of perfect equilibria using LPSs defined over individual strategy sets (see Theorem 2.4). To do so, we introduce a definition of product LPSs of a particular form (see Definition 2.2). This

¹ This approach and interpretation have parallels in Kohlberg and Reny (1997) where a characterization of consistent assessments is obtained using conditional probability systems.

operation seems, initially, to be ad hoc. However, we show in Section 4 that any LPS over strategy profiles that is independent across players has an equivalent LPS that is obtainable using our product formula. Thus, all independent beliefs are included in those obtained from the special form of product LPSs used in our construction.

2. Definitions and statement of the theorem

We consider finite games in normal form. Let $\mathcal{N} = \{1, \dots, N\}$ be the set of players. For each $n \in \mathcal{N}$, let S_n be player n 's set of pure strategies, and let Σ_n be his set of mixed strategies. Denote $S = \prod_{n \in \mathcal{N}} S_n$, and $\Sigma = \prod_{n \in \mathcal{N}} \Sigma_n$. As usual, for each $n \in \mathcal{N}$, we let $S_{-n} = \prod_{m \neq n} S_m$ and $\Sigma_{-n} = \prod_{m \neq n} \Sigma_m$. With these sets of players and strategies fixed, the space of games, call it Γ , is a Euclidean space of dimension $N|S|$ in which each point assigns to each of the players a payoff at each of the $|S|$ profiles of pure strategies. Given a game $G \in \Gamma$, and a mixed-strategy profile $\sigma \in \Sigma$, $G_n(\sigma)$ denotes player n 's expected payoff in G from σ . According to Selten (1975) a mixed-strategy profile σ^* is a perfect equilibrium of a game G if there exists a sequence $\{\sigma^k\}_{k=1}^\infty$ of completely mixed strategy profiles (i.e., in the interior of Σ) converging to σ^* , and such that for all k , σ^* is a best reply to σ^k in the game G .

Definition 2.1. Let X be a finite set. A *lexicographic probability system (LPS)* of order K over X is a $(K+1)$ -tuple $r = (r^0, \dots, r^K)$ of probability distributions on X . An LPS r is said to have *full support* if $\bigcup_{k=0}^K \text{supp } r^k = X$.

For each player n , we use ρ_n to denote an LPS over S_n . ρ_n is interpreted as a collection of theories (held in common by n 's opponents) about n 's strategy choice: ρ_n^0 is the primary theory, ρ_n^1 the secondary theory, ρ_n^2 the tertiary theory, and so on. In particular, it is considered infinitely more likely that player n chooses strategy ρ_n^k than that he chooses ρ_n^{k+1} .

By an LPS profile of order K , we mean an N -tuple $\rho = (\rho_1, \dots, \rho_N)$ where for each n , ρ_n is an LPS of order K over S_n . An LPS profile ρ has full support if each ρ_n has full support; in this case, define $\ell(\rho) = \max_n \min\{k : \bigcup_{i=0}^k \text{supp } \rho_n^i = S_n\}$, the order of ρ required to get full support for all players.

Given an LPS profile, our next objective is to define for each player n an LPS over S_{-n} , interpreted as n 's beliefs about others' actions. The particular construction we use has the following intuitive justification. Consider an LPS profile ρ of order 1 for a 3-player game. For a player, say 1, what should his beliefs over $S_2 \times S_3$ be? He believes that each of his opponents is infinitely more likely to play his first mixed strategy than his second. If we assume that the deviation of either player to his second strategy is independent of the other's deviation and equally likely, then the following LPS, μ_1 , of order 2 would be a reasonable assignment of beliefs for player 1: $\mu_1^0 = \rho_2^0 \times \rho_3^0$; $\mu_1^1 = \frac{1}{2}(\rho_2^0 \times \rho_3^1) + \frac{1}{2}(\rho_2^1 \times \rho_3^0)$; $\mu_1^2 = \rho_2^1 \times \rho_3^1$. (We show in Section 4 that actually the seemingly arbitrary equal weightings by $\frac{1}{2}$ in the formula for μ_1^1 is not restrictive.) The definition below generalizes this notion to N -player games and LPS profiles of any finite order.

Definition 2.2. Given an LPS profile ρ of order K , the *induced lexicographic beliefs for player n* is the product LPS μ_n of order $K(N-1)$ over S_{-n} obtained

as follows. For $k = 0, \dots, K(N - 1)$,

$$\mu_n^k = \sum_{\substack{(k_1, \dots, k_{n-1}, k_{n+1}, \dots, k_N) \\ k_1 + \dots + k_{n-1} + k_{n+1} + \dots + k_N = k}} C^k (\rho_1^{k_1} \times \dots \times \rho_{n-1}^{k_{n-1}} \times \rho_{n+1}^{k_{n+1}} \times \dots \times \rho_N^{k_N}),$$

where C^k is the appropriate normalizing constant.²

We will now illustrate the idea of lexicographic beliefs using a simple example. Suppose we have a three-player game where player 1’s strategy set is $S_1 = \{A, B, C\}$ and player 2’s strategy set is $S_2 = \{a, b, c\}$. Consider an LPS profile of order three in which the LPSs of player 1 and 2 are as follows:

ρ^0	1	—	—	and	ρ^0	1	—	—
ρ^1	—	$\frac{1}{2}$	$\frac{1}{2}$		ρ^1	—	$\frac{1}{2}$	$\frac{1}{2}$
ρ^2	—	1	—		ρ^2	$\frac{1}{2}$	$\frac{1}{2}$	—
ρ^3	—	—	1		ρ^3	—	—	1

(2.1)

where “—” replaces “0” for readability. The induced lexicographic beliefs for player 3 is the product LPS of order 6 over $S_1 \times S_2$ that is given by:

	Aa	Ab	Ac	Ba	Bb	Bc	Ca	Cb	Cc
ρ^0	1	—	—	—	—	—	—	—	—
ρ^1	—	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{4}$	—	—	$\frac{1}{4}$	—	—
ρ^2	$\frac{1}{6}$	$\frac{1}{6}$	—	$\frac{1}{3}$	$\frac{1}{12}$	$\frac{1}{12}$	—	$\frac{1}{12}$	$\frac{1}{12}$
ρ^3	—	—	$\frac{1}{4}$	$\frac{1}{16}$	$\frac{3}{16}$	$\frac{1}{8}$	$\frac{5}{16}$	$\frac{1}{16}$	—
ρ^4	—	—	—	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{1}{6}$	—	$\frac{1}{6}$	$\frac{1}{3}$
ρ^5	—	—	—	—	—	$\frac{1}{2}$	$\frac{1}{4}$	$\frac{1}{4}$	—
ρ^6	—	—	—	—	—	—	—	—	1

(2.2)

Given a lexicographic belief system, the reverse operation of checking whether it is a product LPS is a straightforward exercise. It involves solving a finite sequence of systems of linear equations. To begin, the 0-th order beliefs must be a product distribution. If they are, then the corresponding mixed strategies describe the 0-th level of the LPS profile. The first level of the LPS profile is now obtained by solving a system of linear equations, after which we can obtain the second level, and so on. It is easy to show that there is a 1-1 correspondence between LPS profiles and the induced lexicographic beliefs.

We now define the optimality of a player’s strategy against an LPS profile. As might be expected, a strategy is optimal against an LPS profile if it is lexicographically optimal against the induced lexicographic beliefs.

Definition 2.3. Let ρ be an LPS profile of order K , and let μ_n be the lexico-

² The assumption that the players’ deviations from their primary strategies are mutually independent yields beliefs that are correlated. For instance, in the 3-player example above, μ_1^1 is an average of two product distributions over $S_2 \times S_3$.

graphic beliefs induced by ρ for player n . A strategy $\sigma_n \in \Sigma_n$ is a *lexicographic best reply of order k to the LPS profile ρ* in a game G if

$$(G_n(\sigma_n, \mu_n^0), \dots, G_n(\sigma_n, \mu_n^k)) \geq_L (G_n(\sigma'_n, \mu_n^0), \dots, G_n(\sigma'_n, \mu_n^k))$$

for all $\sigma'_n \in \Sigma_n$, where \geq_L is the lexicographic ordering on vectors.

We now state our main result.

Theorem 2.4. *There exist positive integers l, K , where $l \leq K$ and l, K depend only on the cardinalities of the strategy sets, such that the following statements are equivalent:*

- (1) σ^* is a perfect equilibrium of a game G .
- (2) there is an LPS profile ρ of order K such that:
 - (α) ρ has full support, with $\ell(\rho) \leq l$;
 - (β) $\rho_n^0 = \sigma_n^*$ for each player n ;
 - (γ) for each player n , σ_n^* is a lexicographic best reply of order K against ρ in G .

3. Proof of the theorem

In this section we prove Theorem 2.4 using three claims. The first claim establishes an equivalence between LPSs and certain polynomial functions of one variable; this result is crucial, both for the proof of Theorem 2.4 and for our characterization of independence in Section 4. The other two claims are technical and their proofs are in the Appendix.

We begin with some definitions concerning polynomials. For a polynomial (or more generally a power series) $f(t) = a_0 + a_1 t^1 + a_2 t^2 + \dots$ in a single variable t , the order of f , denoted $o(f)$, is the smallest integer i for which $a_i \neq 0$. (The order of the zero function is ∞ .) We say that $f > 0$ if $a_{o(f)} > 0$.

Let $\kappa = \sum_{n \in \mathcal{N}} |S_n|$. Given a polynomial function $\eta : \mathbb{R} \rightarrow \mathbb{R}^\kappa$, for each player $n \in \mathcal{N}$ and pure strategy $s_n \in S_n$ we will write $\eta_{s_n} = \sum_i \eta_{s_n, i} t^i$ for the corresponding coordinate function. The order of η is given by $o(\eta) \equiv \max_{n, s_n} o(\eta_{s_n})$. For every game G , and every player n , the payoff function G_n has an obvious extension from strategy profiles to polynomial functions η , given by $G_n(\eta) = \sum_s G_n(s) \prod_m \eta_{s_m}$. We say that a strategy τ_n for player n is a best reply of order k against a polynomial η in a game G if for all $\tilde{s}_n \in S_n$, $G_n(\tilde{s}_n, \eta_{-n}) - G_n(\tau_n, \eta_{-n})$ is either nonpositive or of order at least $k + 1$. If $k = \infty$, we say that τ_n is a best reply against η .

Claim 3.1. *Statement 2 of Theorem 2.4 is equivalent to the existence of a polynomial function $\eta : \mathbb{R} \rightarrow \mathbb{R}^\kappa$ such that*

- (a) $\eta_{s_n} > 0$ for all n, s_n , with $o(\eta) \leq l$;
- (b) $\eta(0) = \sigma^*$;
- (c) for each player n , σ_n^* is a best reply of order K against η in G .

Proof: Given an LPS profile ρ satisfying the conditions of Statement 2 of Theorem 2.4, the polynomial function defined by $\eta_n = \sum_{i=0}^K t^i \rho_n^i$ satisfies the conditions of the Claim. So it remains to establish the sufficiency of these conditions.

Given a polynomial η satisfying conditions (a) to (c) of the Claim, observe first that every function of the form $(\sum_i \eta_{s_n,i}(ct)^i)_{n,s_n}$ (for $c > 0$) or $((1 + \sum_{i>0} c_{n,i}t^i)\eta_n)_{n \in \mathcal{N}}$, satisfies (a) to (c). Indeed, for $c > 0$, η is obtained from the function $(\sum_i \eta_{s_n,i}(ct)^i)_{n,s_n}$ by a positive transformation of the variable t . Hence, the latter satisfies conditions (a) to (c) iff η does. To prove that the second type of transformation also produces a polynomial with properties (a) to (c), consider a function $\zeta = ((1 + \sum_{i>0} c_{n,i}t^i)\eta_n)_{n \in \mathcal{N}}$. Obviously, $\zeta(0) = \eta(0) = \sigma^*$, and ζ satisfies (b). For each n , $1 + \sum_{i>0} c_{n,i}t^i$ is a positive polynomial with order 0; since, for each n, s_n , η_{s_n} is positive and has order no more than l , ζ_{s_n} is also positive and $o(\zeta_{s_n}) = o(\eta_{s_n}) \leq \ell$. Thus ζ satisfies (a). Finally, for each n, s_n , $G_n(s_n, \zeta_{-n}) - G_n(\sigma_n^*, \zeta_{-n}) = \prod_{m \neq n} (1 + \sum_{i>0} c_{m,i}t^i)(G_n(s_n, \eta_{-n}) - G_n(\sigma_n^*, \eta_{-n}))$. Therefore $G_n(s_n, \zeta_{-n}) - G_n(\sigma_n^*, \zeta_{-n})$ and $G_n(s_n, \eta_{-n}) - G_n(\sigma_n^*, \eta_{-n})$ have the same sign and the same order. Thus, ζ satisfies (c) as well.

We use a sequence of three transformations of the sort given in the previous paragraph to obtain a function ζ such that for each n , and $0 \leq k \leq K$, the vector $(\zeta_{s_n,k})_{s_n \in S_n}$ is a probability distribution. The result then follows immediately.

The first transformation ensures that for all n, s_n and $0 \leq i \leq K$, $\eta_{s_n,i} \geq 0$ and $\bar{\eta}_{n,i} \equiv \sum_{s_n \in S_n} \eta_{s_n,i} > 0$. The construction is done inductively. By (b), this condition is obviously true for $i = 0$. Assume now that the function η is such that for each n, s_n , and $i < k$, $\eta_{s_n,i} \geq 0$ and $\bar{\eta}_{n,i} > 0$. If for some n, s_n , we have that $\eta_{s_n,k} < 0$, then by (a), there exists $i < k$ such that $\eta_{s_n,i} > 0$. Therefore, there exists a constant $c > 0$ such that for the function $(1 + ct + \dots + ct^k)\eta$, the coefficients of t^k are all nonnegative and $\bar{\eta}_{n,k} > 0$ for all n .

The next transformation gives us that $\sum_{i=1}^K \bar{\eta}_{n,i} < 1$ for all n . To obtain this, choose $c > 0$ sufficiently small and consider the function $\zeta(t) = \eta(ct)$.

Finally, we choose positive constants $c_{n,i}$ for $n \in \mathcal{N}$, $1 \leq i \leq K$ such that for the function $\zeta_n = (1 + \sum_{i=1}^K c_{n,i}t^i)\eta_n$, the vector of coefficients $(\zeta_{s_n,k})_{s_n \in S_n}$ is a probability distribution for all $1 \leq k \leq K$. (The constants in ζ_n equal those in η_n and therefore yield a probability distribution.) Observe that the constants have to satisfy the condition:

$$\left(1 + \sum_{i=1}^K c_{n,i}t^i\right) \left(1 + \sum_{j>0} \bar{\eta}_{n,j}t^j\right) = 1 + t + \dots + t^K + \dots$$

The choice of constants is made inductively, as follows. Define $c_{n,1} = 1 - \bar{\eta}_{n,1}$. We have that $0 < c_{n,1} < 1$ for all n , as $0 < \bar{\eta}_{n,1} \leq \sum_{i=1}^K \bar{\eta}_{n,k} < 1$. Now assume that we have chosen constants $0 < c_{n,i} < 1$ for all n and $1 \leq i \leq k$. We can choose $c_{n,k+1}$ to be the number $1 - (c_{n,k}\bar{\eta}_{n,1} + \dots + c_{n,1}\bar{\eta}_{n,k} + \bar{\eta}_{n,k+1})$. Since the $\bar{\eta}_{n,i}$'s, the $c_{n,i}$'s and $\sum_{i=1}^K \bar{\eta}_{n,i}$ are strictly between 0 and 1, $0 < c_{n,k}\bar{\eta}_{n,1} + \dots + c_{n,1}\bar{\eta}_{n,k} + \bar{\eta}_{n,k+1} \leq \sum_{i=1}^K \bar{\eta}_{n,i} < 1$. Therefore, $c_{n,k+1}$ is also between 0 and 1. With our choice of the constants $c_{n,i}$, the coefficients of ζ give us the LPS required by Statement 2. □

Suppose that a strategy profile σ is not a perfect equilibrium of a game G . Then there exists a closed neighbourhood U of σ in Σ such that σ is not a best reply against any completely mixed strategy in U . For each $\tau \in U$, let $g(\tau) = \min_{n,s_n} \tau_{s_n}$ and let $f(\tau) = \max_{n,s_n} G_n(\tau_{-n}, s_n) - G_n(\tau)$. Then, $f^{-1}(0) \subseteq g^{-1}(0)$ and Łojasiewicz's Inequality (cf. Lemma A.1 in the Appendix) implies that there exist an integer m_σ and a constant c such that $g^{m_\sigma} \leq cf$. If η were a

strictly positive polynomial such that $\eta(0) = \sigma$, then for all sufficiently small $t > 0$, $\eta(t) \in U$ and, by our previous argument, $o(f(\eta)) \leq o(\eta)m_\sigma$. In other words, σ cannot be a best reply of order $o(\eta)m_\sigma$ against η . The following Claim states that we can take this integer m_σ to be independent of the strategy profile σ and the game G .

Claim 3.2. *There exists a positive integer m such that the following condition is sufficient for a strategy profile σ to be a perfect equilibrium of a game G : there exists a polynomial function $\eta \gg 0$ with $\eta(0) = \sigma$ and such that σ is a best reply of order $o(\eta)m$ against η in G .*

Consider a perfect equilibrium σ . By definition, it is a best reply against a sequence of completely mixed strategy profiles converging to it. Therefore, by the Nash Curve Selection Lemma (Cf. Proposition 8.1.13 of Bochnak, Coste, and Roy, 1998) there exists an analytic function $\eta : [0, \varepsilon) \rightarrow \Sigma$ such that: (i) η is strictly positive with order, say, l_σ ; (ii) $\eta(0) = \sigma$; and (iii) σ is a best reply against η . Claim 3.3 below asserts the existence of a uniform upper bound on l_σ .

Claim 3.3. *There exists a positive integer l such that the following condition is necessary for a strategy profile σ to be a perfect equilibrium of a game G : there exists an analytic function $\eta : [0, \varepsilon) \rightarrow \mathbb{R}_+^K$ such that $o(\eta) \leq l$, $\eta(0) = \sigma$, and σ is a best reply to $\eta(t)$ for all $t \in [0, \varepsilon)$.*

We are now ready to prove the Theorem. By Claim 3.1, it is sufficient to prove that Statement 1 of Theorem 2.4 is equivalent to the existence of a polynomial η with properties (a) to (c). Let now m and l be as in Claims 3.2 and 3.3, respectively. Define K to be lm . Given a strategy profile σ^* , if there exists a polynomial function η satisfying the conditions of Claim 3.1, then by Claim 3.2, it is perfect. To prove the necessity of Statement 2, suppose σ^* is a perfect equilibrium. Then by Claim 3.3, there exists an analytic function $\eta : [0, \varepsilon) \rightarrow \mathbb{R}_+^K$ such that $o(\eta) \leq l$, $\eta(0) = \sigma^*$, and σ^* is a best reply against $\eta(t)$ for all small t . The polynomial obtained by truncating η to its first $K + 1$ terms satisfies (a) to (c) of Claim 3.1.

4. Independence

In this section, we examine the issue of independence for LPSs defined on product spaces. Specifically, we give a definition of independence and show that any LPS that is independent is equivalent to one that is obtained using the product rule of Definition 2.2. Thus, the product formula is not ad hoc, but rather a canonical representation of independent beliefs.

The definition of independence here is a version of what BBD (1991a) call strong independence. However, the non-Archimedean field we use is the ordered field $\mathbb{R}(t)$ of rational functions in one indeterminate.³ The ordering on $\mathbb{R}(t)$ is given by the following. First, as in Section 3, we say that a polynomial $f = \sum_i a_i t^i$ is positive if $a_{o(f)} > 0$, where $o(f)$ is the order of f . Then, a rational function $f(t)/g(t)$ is positive if $f(t)g(t) > 0$. This field was first employed in game theory by Hammond (1994) who argued that it would be an

³ The conclusions of this Section would be valid for the field of Puiseux series (which contains $\mathbb{R}(t)$) but not for the field \mathfrak{R} of hyperreals used by BBD.

appropriate field for studying refinements, since it is in some sense the smallest and simplest ordered field that is both non-Archimedean and an extension of \mathbb{R} . $\mathbb{R}(t)$ truly captures the notion of infinitesimals – for example, the “number” t is positive but smaller than every positive real number.

Definition 4.1. Let X be a finite set, and let $r = (r^0, r^1, \dots, r^K)$ be an LPS of order K over X .

- (1) An $\mathbb{R}(t)$ -valued probability distribution on X is a function $P : X \rightarrow \mathbb{R}(t)$ such that $P(x) \geq 0$ for all x and $\sum_{x \in X} P(x) = 1$.
- (2) An $\mathbb{R}(t)$ -valued probability distribution P on X is equivalent to r if there exist positive polynomials f^0, f^1, \dots, f^K such that $o(f^0) < o(f^1) < \dots < o(f^K)$ and

$$P = \left(\sum_{i=0}^K f^i \right)^{-1} \left(\sum_{j=0}^K f^j r^j \right).$$

The reason for deeming P equivalent to r in the definition above is purely decision-theoretic.⁴ BBD (1991a, Theorem 3.1) show that in an Anscombe-Aumann type subjective expected utility framework with a weakened Archimedean axiom, one obtains an LPS over the state space, instead of a unique probability distribution. An equivalent $\mathbb{R}(t)$ -valued distribution P can also be interpreted as representing an agent’s subjective beliefs, in the sense that if we allow $\mathbb{R}(t)$ -valued utility functions, this agent’s preferences can be represented using a utility function that involves taking expectations (of a real-valued utility function) w.r.t. P – see BBD, 1991a, Theorem 6.1.

Let $S = \prod_n S_n$ be a finite state space. (For example, S is the set of all pure strategy profiles in a game.) Let $\mu = (\mu^0, \mu^1, \dots, \mu^K)$ be an LPS on S .

Definition 4.2. μ is an independent LPS if there exists an equivalent $\mathbb{R}(t)$ -valued probability distribution on S that is a product distribution.

Observe that the product formula given in Definition 2.2 induces an independent LPS on strategy profiles. Indeed, let $\rho = (\rho_1, \dots, \rho_N)$ be an LPS profile of order K . For each n , let $P_n = (\sum_{i=0}^K t^i)^{-1} (\sum_{k=0}^K t^k \rho_n^k)$. Then, for each player n , $\prod_{m \neq n} P_m$ is equivalent to his beliefs μ_n over S_{-n} . More generally, given positive coefficients $a_{n,i}$ for each n and $0 \leq i \leq K$, the LPS μ_n over S_{-n} given by

$$\mu_n^k = \sum_{\substack{(k_1, \dots, k_{n-1}, k_{n+1}, \dots, k_N) \\ k_1 + \dots + k_{n-1} + k_{n+1} + \dots + k_N = k}} C^k \left(\prod_{m \neq n} a_{m, k_m} \rho_m^{k_m} \right)$$

where C^k is the appropriate normalizing constant, is an independent LPS.

Our notion of independence is closely related to BBD’s notion of strong independence for LPSs (see BBD, p. 90). μ is strongly independent if there exists an equivalent \mathfrak{R} -valued distribution that is a product distribution. (Here \mathfrak{R} is the space of hyperreals.) Equivalently, μ is strongly independent if there exists a sequence $\lambda(n)$ in $(0, 1)^K$ converging to zero such that for each

⁴ The notion of equivalence here (and in Definition 4.3) is not the standard mathematical one, since it is not symmetric.

n , the nested convex combination $(1 - \lambda_1(n))\mu^0 + \lambda_1(n)[(1 - \lambda_2(n))\mu^1 + \lambda_2(n) \cdot [\dots + \lambda_{(K-1)}(n)[(1 - \lambda_K(n))\mu^{K-1} + \lambda_K(n)\mu^K]]]$ is a product distribution. Since $\mathbb{R}(t) \subset \mathfrak{R}$, we have that an LPS μ that is independent under our definition is also strongly independent. However, the converse is not true. For example consider distributions on each S_n that have values in the ring of analytic functions, at least one of which is not a rational function. The resulting product is clearly strongly independent but not independent in the sense of Definition 4.2.

Definition 4.3. Let μ be an LPS of order K over S . We say that an LPS ν of order L is equivalent to μ if there exists a monotonic function $\pi : \{0, 1, \dots, K\} \rightarrow \{0, 1, \dots, L\}$ such that for each $0 \leq l \leq L$, $\nu^l = \sum_{k:\pi(k) \leq l} a_k \mu^k$ for some set of constants a_k with the property that $a_k > 0$ if $\pi(k) = l$.

As with $\mathbb{R}(t)$ -valued distributions, this equivalence stems from decision-theoretic considerations: BBD (1991a, Theorem 3.1) identify this class of LPSs as the set of distributions that arise from preferences. The following Proposition gives a characterization of independent LPSs in terms of our product formula.

Proposition 4.4. Let μ be an independent LPS of order K over S . Then there exists an equivalent LPS ν of order $L \geq K$ and an LPS profile $\rho = (\rho_1, \dots, \rho_N)$ of order L such that ν is obtained from ρ using the formula in Definition 2.2.

Proof: We can assume without loss of generality that μ has full support. Since μ is independent, there exist positive polynomials f^0, f^1, \dots, f^K , with $o(f^0) < o(f^1) < \dots < o(f^K)$, and for each n, s_n , a positive polynomial $\eta_{s_n} = \sum_{i \geq 0} \eta_{s_n, i} t^i$ such that for each $s = (s_1, \dots, s_N) \in S$, $\sum_{j=0}^K f^j \mu^j(s) = \prod_n \eta_{s_n}$. We can choose the polynomials such that f^0 and the η_n 's have 1 as their constant term. (In particular, all of them have order zero.) Let $L = o(f^K)$. Clearly, $o(\eta) \leq L$.

By a quasi LPS of order L , we mean a collection $\nu = (\nu^0, \nu^1, \dots, \nu^L)$ where $\nu^i \in \mathbb{R}^{|S|}$ for all i . From the polynomial η , form now a quasi LPS ν_η as follows: for each $s = (s_1, \dots, s_N) \in S$, and $0 \leq i \leq L$, let $\nu_\eta^i(s)$ be the coefficient of t^i in $\prod_n \eta_{s_n}$ (or equivalently the coefficient of t^i in $\sum_{j=0}^K f^j \mu^j(s)$). Observe that, modulo a normalization, the formula for obtaining ν_η from η is the same as the product formula in Definition 2.2. Let $\pi : \{0, 1, \dots, K\} \rightarrow \{0, 1, \dots, L\}$ be the function $\pi(j) = o(f^j)$. Then each vector ν_η^i is a linear combination of vectors μ^j , for j such that $\pi(j) \leq i$, that assigns a positive weight to $\pi^{-1}(i)$, if it exists, since the f^j 's are positive. Thus ν_η is “equivalent” to μ in the sense of Definition 4.3.

We claim now that the two types of transformations that were used in the proof of Claim 3.1 induce quasi LPSs that are equivalent to μ . Consider first the transformation $\zeta(t) = \eta(ct)$ ($c > 0$). The quasi LPS ν_ζ obtained from ζ has the property that for each l , $\nu_\zeta^l = c^l \nu_\eta^l$. Hence ν_ζ is equivalent to μ . For the transformation $\varsigma_n = (1 + \sum_{i>0} c_{n,i} t^i) \eta_n$, we have for each l ,

$$\nu_\varsigma^l = \sum_{i_1 + \dots + i_N + j = l} \left(\prod_n c_{n, i_n} \right) \nu_\eta^j,$$

where $c_{n,0} = 1$ for all n . Therefore, ν_ς is also equivalent to μ .

Apply the transformations of Claim 3.1 to obtain a polynomial ζ such that for each n , and $0 \leq i \leq L$, the coefficients $\zeta_{s_n, i}$ form a probability distribution over S_n . The product formula gives an LPS v of order L which, by the previous arguments, is equivalent to μ . \square

There is an odd aspect to Proposition 4.4, in that the LPS profile ρ – from which the LPS v is obtained using the product formula – has the same order as v .⁵ Applying the product formula to ρ , though, yields an LPS of order LN . The question then arises as to whether this larger LPS is also equivalent to μ . The answer, in general, is no. If, however, none of the polynomials f^j has a term with a negative coefficient, then one can obtain the following sharper statement: there exists an LPS profile ρ of order $\ell \geq K/N$ such that the corresponding LPS v over S of order ℓN is equivalent to μ .

The above observations suggest that we should perhaps strengthen the conditions for an $\mathbb{R}(t)$ -valued probability distribution to be equivalent to an LPS by requiring that the polynomials f^j not contain any negative terms.⁶ From the view point of decision theory, there seems to be no justification for such a strengthening.

5. Extensions and an example

Since extensive-form perfect equilibria are the perfect equilibria of the agent-normal form, our result extends to a characterization of the former as well. When we apply our methods to sequential equilibria (Kreps and Wilson, 1982) we get precisely the characterization of sequential equilibria in terms of conditional probability systems as in McLennan (1989). To see this, take an LPS profile of behavioural strategies for an extensive-form game. Now compute the induced lexicographic beliefs at each information set of each player, using the product operation given in Definition 2.2. The 0-th order beliefs of the system yield the relevant components of the conditional probability system associated with the LPS profile.⁷ Sequential rationality requires optimality of order zero at each information set given the lexicographic beliefs and the equilibrium strategies of the opponents. Since only the 0-th order beliefs matter, the LPS profile can be taken such that various levels have disjoint supports.

It can also be shown that proper equilibria (Myerson, 1978) have a characterization of the form given in Theorem 2.4. The only additional constraint is that if a pure strategy s_n for player n does worse (lexicographically) than another pure strategy \tilde{s}_n (for the same player) against the LPS profile ρ , then $\min\{k : \rho_n^k(s_n) > 0\} > \min\{k : \rho_n^k(\tilde{s}_n) > 0\}$.

Theorem 2.4 is not the most economical characterization possible. To see this, remark that equilibria in completely mixed strategies are always perfect; and yet, we require even those to be supported by higher order beliefs. To

⁵ This feature shows up in Theorem 2.4 as well, where ρ has order K and the equilibrium is required to be a best reply only of order K – and not $K(N - 1)$ – against ρ .

⁶ Hammond (1994) considers precisely such functions, which he calls rational probability functions.

⁷ Because beliefs are obtained using the product formula, we obtain the necessary condition for consistent assessments given by Kreps and Ramey (1987), namely that beliefs at each information set are obtained by averaging over a finite number of product distributions induced by behavioural strategy profiles.

take care of this redundancy in beliefs, consider the following variation of our result. Let m and l be as in Claims 3.2 and 3.3, respectively. A strategy profile σ is a perfect equilibrium of a game G if there exists an LPS profile ρ of order $L \leq lm$ such that: (x) ρ has full support, with $\ell(\rho) \leq l$; (y) $\rho^0 = \sigma$; and (z) σ is a best reply of order $\min(L(N - 1), \ell(\rho)m)$ against ρ . It was for the sake of conceptual clarity that we chose to state the Theorem the way we did.

Our characterization is only essentially finite. The fact that we do not have an explicit formula for l or K means that we are still faced with an arbitrarily high number of computations in practice. We hope that future research can provide tight bounds on K and l .

We conclude with an example comparing the LPSs obtained in BBD and the ones used here. For simplicity in presentation, we merely look at a sequence of mixed strategies without any reference to perfection. Suppose there are two players, 1 and 2, with $S_1 = \{A, B, C\}$ and $S_2 = \{a, b, c\}$. Assume that (A, a) is a perfect equilibrium and parameterize the probability with which each pure strategy is played in a sequence of trembles by rational functions of t :

$$\frac{1}{1 + t + t^2 + t^3} \{1, .5t + t^2, .5t + t^3\} \quad \text{and} \\ \frac{1}{1 + t + t^2 + t^3} \{1 + .5t^2, .5t + .5t^2, .5t + t^3\}. \tag{5.1}$$

Using the model developed here, the LPS profile we get is given by Equation 2.1 in Section 2 and the induced product LPS over $S_1 \times S_2$ is given by Equation 2.2.

By contrast, applying BBD’s algorithm (see proof of Proposition 2 in BBD), we derive the following LPS over the joint action space:

	Aa	Ab	Ac	Ba	Bb	Bc	Ca	Cb	Cc
ρ^0	1	—	—	—	—	—	—	—	—
ρ^1	—	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{4}$	—	—	$\frac{1}{4}$	—	—
ρ^2	—	$\frac{1}{5}$	—	$\frac{2}{5}$	$\frac{1}{10}$	$\frac{1}{10}$	—	$\frac{1}{10}$	$\frac{1}{10}$
ρ^3	—	—	—	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{5}$	$\frac{1}{20}$	$\frac{3}{20}$	$\frac{1}{10}$
ρ^4	—	—	—	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{1}{6}$	—	$\frac{1}{6}$	$\frac{1}{3}$
ρ^5	—	—	—	—	—	$\frac{3}{14}$	—	$\frac{3}{14}$	$\frac{4}{7}$
ρ^6	—	—	—	—	—	—	—	—	1

(5.2)

The induced marginal LPSs over the individual strategy sets are given by:

	A	B	C		a	b	c
ρ^0	1	—	—	and	ρ^0	1	—
ρ^1	$\frac{1}{2}$	$\frac{1}{4}$	$\frac{1}{4}$		ρ^1	$\frac{1}{2}$	$\frac{1}{4}$
ρ^2	$\frac{1}{5}$	$\frac{3}{5}$	$\frac{1}{5}$		ρ^2	$\frac{2}{5}$	$\frac{2}{5}$
ρ^3	—	$\frac{7}{10}$	$\frac{3}{10}$		ρ^3	$\frac{3}{10}$	$\frac{4}{10}$
ρ^4	—	$\frac{1}{2}$	$\frac{1}{2}$		ρ^4	$\frac{1}{6}$	$\frac{1}{3}$
ρ^5	—	$\frac{3}{14}$	$\frac{11}{14}$		ρ^5	—	$\frac{3}{14}$
ρ^6	—	—	1		ρ^6	—	1

(5.3)

Appendix A

This appendix contains proofs of Claims 3.2 and 3.3. We begin by stating (without proof) a useful Lemma from semi-algebraic geometry – cf. Bochnak, Coste, and Roy, 1998, Corollary 2.6.7.

Lemma A.1. (*Łojasiewicz's Inequality*) *Let $X \subset \mathbb{R}^m$ be a compact, semi-algebraic set. Suppose $f : X \rightarrow \mathbb{R}$ and $g : X \rightarrow \mathbb{R}$ are continuous semi-algebraic functions with $f^{-1}(0) \subseteq g^{-1}(0)$. Then there exists a constant $c \in \mathbb{R}$ and an integer $k > 0$ such that $|g(x)|^k \leq c|f(x)|$ for all $x \in X$.*

Our proof technique exploits the semi-algebraic structure of the graphs of the equilibrium correspondence and the best reply correspondence. Therefore, we first develop some machinery. Remark that Claims 3.2 and 3.3 depend on the game G only through its best reply correspondence. Since best reply correspondences are preserved under positive affine transformations of the payoff functions, there is no loss of generality in assuming that the space of games Γ is the unit cube in $\mathbb{R}^{N|S|}$. Let E be the graph of the Nash equilibrium correspondence, i.e.,

$$E = \{(G, \sigma) \in \Gamma \times \Sigma \mid \sigma \text{ is a Nash equilibrium of } G\}.$$

Let $P \subset E$ be the graph of the perfect equilibrium correspondence. Both E and P are semi-algebraic subsets of $\Gamma \times \Sigma$. (See, for instance, Blume and Zame, 1994.) Furthermore, E is compact.

For $\theta \in \mathbb{R}^k$, and $n \in \mathcal{N}$, let $\bar{\theta}_n = \sum_{s_n \in S_n} \theta_{s_n}$. Given $\sigma \in \Sigma$, and $\theta \in \mathbb{R}^k$ with $\bar{\theta}_n > -1$ for all n , define $\sigma(\theta) = \left(\frac{\sigma_n + \theta_n}{1 + \bar{\theta}_n} \right)_n$. Observe that a strategy profile σ is a perfect equilibrium of a game G iff there exists a sequence of θ 's in \mathbb{R}_{++}^k converging to zero such that σ is a best reply against $\sigma(\theta)$ all along the sequence. Let

$$X = \{((G, \sigma), \theta) \in E \times [-1, 1]^k \mid \bar{\theta}_n \geq -1/2 \forall n, \text{ and } \sigma(\theta) \in \Sigma\},$$

and let $\partial X = \{((G, \sigma), \theta) \in X \mid \sigma(\theta) \in \partial \Sigma\}$, where $\partial \Sigma$ is the topological boundary of Σ in the affine space generated by Σ . X is easily verified to be a compact, semi-algebraic set. Also, $E \times \{0\} \subset X$. Denote by p the natural projection from X onto E . For each n, s_n , let

$$X_{s_n}^0 = \{(G, \sigma, \theta) \in X \mid \theta_{s_n} = 0\},$$

and

$$X_{s_n}^1 = \{(G, \sigma, \theta) \in X \mid s_n \text{ is a best reply to } \sigma(\theta) \text{ in } G\}.$$

By the Generic Local Triviality Theorem (cf. Bochnak, Coste, and Roy, 1998, Theorem 9.3.2) there exist: (i) a partition of E into a finite number of semi-algebraic subsets E^1, \dots, E^k ; (ii) for each $1 \leq i \leq k$, a triple of semi-algebraic fibres $(F^i, \partial F^i, \{f_0^i\})$; and (iii) for each i , a homeomorphism $h^i : E^i \times (F^i, \partial F^i, \{f_0^i\}) \rightarrow (p^{-1}(E^i), p^{-1}(E^i) \cap \partial X, E^i \times \{0\})$ such that $p \circ h^i((G, \sigma), f) = (G, \sigma)$ for all $((G, \sigma), f) \in E^i \times F^i$. Moreover, the trivialization can be made compatible with each $X_{s_n}^j$ for n, s_n and $j = 0, 1$, in the

sense that for $1 \leq i \leq k$, there exists a semi-algebraic set $F_{s_n}^{ij} \subset F^i$ such that $h^i(E^i \times F_{s_n}^{ij}) = p^{-1}(E^i) \cap X_{s_n}^j$.

Triangulate E such that (i) both P and the sets E^i are all unions of interiors of simplices, and (ii) the support function is constant over the interior of each simplex, i.e., for any two points (G, σ) and (G', σ') belonging to the interior of a simplex τ , $\text{supp}(\sigma_n) = \text{supp}(\sigma'_n) \equiv S_n(\tau)$ for all n . Call the space of this triangulation \mathcal{E} . We use d to denote the distance function on E that is induced by the triangulation. For each i , triangulate the fibre F^i such that both ∂F^i and the sets $F_{s_n}^{ij}$ are all subcomplexes, and f_0^i is a vertex. Take still a barycentric subdivision, so that the subcomplexes of the original triangulation are now full. (Thus the intersection of any one of these subcomplexes with a simplex of the refined triangulation is a face of the simplex.) Call the space of this triangulation \mathcal{F}^i . Finally, given a simplex τ , we use τ^0 to denote its interior.

Proof of Claim 3.2: Let Q be the (nonempty) collection of simplices τ of \mathcal{E} such that $\tau^0 \not\subseteq P$. By virtue of our triangulation, $\tau \in Q$ implies $\tau^0 \cap P = \emptyset$. Fix now $\tau \in Q$. There exists a unique i such that $\tau^0 \subseteq E^i$. Let \mathcal{F}_0^i be the (nonempty) set of simplices of \mathcal{F}^i that have $\{f_0^i\}$ as a vertex and are not contained in ∂F^i . Choose $v \in \mathcal{F}_0^i$. Let $v_0 = v \cap \partial F^i$. v_0 is then a proper face of v . Since the equilibria in τ^0 are not perfect, $(\bigcap_{n \in \mathcal{N}, s_n \in S_n(\tau)} F_{s_n}^{i1}) \subseteq v_0$. To see this, remark that the intersection of each of the sets $F_{s_n}^{i1}$ with v is a face of v (that contains f_0^i as a vertex, since points in τ^0 are equilibria with support $S(\tau)$). Thus, if their common intersection was not contained in v_0 , it would contain a vertex $f_1^i \in v \setminus v_0$ and hence also the simplex $[f_0^i, f_1^i]$. The image of $(\tau^0 \times (f_0^i, f_1^i))$ under h^i would then render all equilibria in τ^0 perfect, contradicting the assumption that $\tau^0 \cap P = \emptyset$. Therefore, $(\bigcap_{n \in \mathcal{N}, s_n \in S_n(\tau)} F_{s_n}^{ij}) \subseteq v_0$. For each player n , denote by $u_n(\tau)$ the uniform mixture over $S_n(\tau)$. Let $Y_{\tau, v}$ be the closure of $h^i(\tau^0 \times v)$ in X . Define $\varphi, \psi : Y_{\tau, v} \rightarrow \mathbb{R}_+$ as follows: first, for $((G, \sigma), \theta) \in h^i(\tau^0 \times v)$ let

$$\varphi(G, \sigma, \theta) = d((G, \sigma), \tau \setminus \tau^0) \left(\max_{n, s_n} G_n(s_n, \sigma_{-n}(\theta)) - G_n(u_n(\tau), \sigma_{-n}(\theta)) \right)$$

and

$$\psi(G, \sigma, \theta) = d((G, \sigma), \tau \setminus \tau^0) \left(\min_{n, s_n} \sigma_{n, s_n} + \theta_{n, s_n} \right);$$

next, let φ and ψ be zero everywhere else. It is easily checked that φ and ψ are continuous functions. Moreover, $\varphi^{-1}(0) \subseteq \psi^{-1}(0)$, since $(\bigcap_{n \in \mathcal{N}, s_n \in S_n(\tau)} F_{s_n}^{ij}) \subseteq v_0$. By the Łojasiewicz inequality now, there exists an integer $m(\tau, v) > 0$ and a constant $c(\tau, v) > 0$ such that $c(\tau, \mu)\varphi \geq \psi^{m(\tau, v)}$. Let m (resp. c) be the maximum over all $\tau \in Q$ and $v \in \mathcal{F}_0^i$ of $m(\tau, v)$ (resp. $c(\tau, v)$).

Let $\eta \gg 0$ be a polynomial function such that $\eta(0) = \sigma$. We will show that if σ is not a perfect equilibrium of a game G , then it cannot be a best reply of order $o(\eta)m$ against η . Observe that this statement is trivially true if σ is not a Nash equilibrium, as σ would not even be a best reply of order zero then. Therefore, let us assume that σ is a Nash, but not a perfect, equilibrium. Then (G, σ) belongs to the interior of some $\tau \in Q$. Define $\theta(t) = \eta(t) - \sigma$. For sufficiently small $t > 0$, $((G, \sigma), \theta(t))$ belongs to a set $Y_{\tau, v}$ of the sort con-

structed above. Therefore $c\varphi(G, \sigma, \theta(t)) \geq \psi(G, \sigma, \theta(t))^m$. Since the order of $(\psi(G, \sigma, \theta(t)))$ (viewed as a polynomial in t) is $o(\eta)$, σ cannot be a best reply of order $o(\eta)m$. □

Proof of Claim 3.3: Let \mathcal{P} be the set of simplices of \mathcal{E} whose interiors belong to P . Fix $\tau \in \mathcal{P}$. There exists i such that τ^0 belongs to E^i . Choose now a point $(G^0, \sigma^0) \in \tau^0$. Since σ^0 is a perfect equilibrium of G^0 , there exists a sequence θ^k in $(0, 1)^K$ converging to zero such that σ^0 is a best reply against $\sigma^0(\theta^k)$ (for all k) in the game G^0 . By passing to a subsequence, if necessary, there exists a simplex $v \in \mathcal{F}^i$ such that $(G^0, \sigma^0, \theta^k) \in h^i((G^0, \sigma^0), v^0)$ for all k . Obviously, $F_{s_n}^{i1} \cap v = v$ for all n and $s_n \in S_n(\tau)$. And, f_0^i is one of the vertices of v . Moreover, $v \cap F_{s_n}^{i0}$ is a proper face of v for all n, s_n . Since $\tau^0 \times v^0$ is connected, this last fact implies that for all $(G, \sigma, \theta) \in h^i(\tau^0 \times v^0)$, $\theta \in \mathbb{R}_{++}^K$. Pick now a point f_1^i in v^0 . By the above mentioned properties of v , the line segment μ joining f_0^i and f_1^i in the simplex v is such that for each $(G, \sigma, \theta) \in h^i(\tau^0 \times (\mu \setminus \{f_0^i\}))$: (i) $\theta_{n, s_n} > 0$ for all n, s_n (thus $\sigma(\theta) \in \Sigma \setminus \partial\Sigma$); and (ii) σ is a best reply against $\sigma(\theta)$ in the game G . For simplicity, we will identify the line segment μ with $[0, 1]$ and denote a typical element of this set by x .

Define $\lambda : \tau \times [0, 1] \rightarrow \mathbb{R}_{++}^K$ as follows: first, for $(G, \sigma, x) \in \tau^0 \times [0, 1]$, let $\lambda(G, \sigma, x)$ be the $d((G, \sigma), \tau \setminus \tau^0)\theta$ where θ is such that $(G, \sigma, \theta) = h^i(G, \sigma, x)$; elsewhere, let λ be zero. λ is continuous. Also, for each n, s_n , $\lambda_{n, s_n}^{-1}(0) = (\tau \times \{0\}) \cup ((\tau \setminus \tau^0) \times [0, 1])$. Therefore, by the Łowasiewicz Inequality, there exists a positive integer q and a constant c such that for all n, s_n , $c\lambda_{s_n}(G, \sigma, x) \geq (x^q)(d((G, \sigma), \tau \setminus \tau^0))^q$.

For each n, s_n , let A_{s_n} be the set of $((G, \sigma), x, y) \in \tau^0 \times (0, 1) \times \mathbb{R}_{++}$ for which there exists θ such that $h^i(G, \sigma, x) = \theta$ and $\theta_{s_n} = y$. $A_{n, s}$ is a lower dimensional semi-algebraic set of $\Gamma \times \Sigma \times \mathbb{R}^2$. Hence there exists a finite integer J such that A_{s_n} is a finite union over $1 \leq j \leq J$ of sets of the form $\{(G, \sigma, x, y) \mid F_{s_n}^j(\cdot) = 0, g^1(\cdot) > 0, \dots, g^{M_j}(\cdot) > 0\}$, where $F_{s_n}^j$ and the g 's are polynomials. Since for each $(G, \sigma) \in \tau^0$, the set of (x, y) s.t. $(G, \sigma, x, y) \in A_{n, s}$ is a 1-dimensional set, the total degree in x and y of each $F_{s_n}^j$ is nonzero. Let d_{s_n} be the maximum over j of this total degree.

By the Nash Curve Selection Lemma (loc. cit) there exist, for each $(G, \sigma) \in \tau^0$ and for each n, s_n , analytic functions $x, y : (-\varepsilon, \varepsilon) \rightarrow \mathbb{R}$ such that $x(0) = y(0) = 0$, and $(G, \sigma, x(t), y(t)) \in A_{s_n}$ for all $t > 0$ (in particular there exists j s.t. $F_{s_n}^j(G, \sigma, x(t), y(t))$ is zero for all t). We can also assume that $x = t^r$ for some positive integer $r \leq d_{s_n}$, because $x(t)$ is nonzero and the total degree of $F_{s_n}^j$ in x and y is no more than d_{s_n} . Therefore, by the Łowasiewicz Inequality, the order of y is no more than qd_{s_n} .

Let $l(\tau) = (\times_{n, s_n} d_{s_n})q$. By the previous paragraph, there exists for each $(G, \sigma) \in \tau^0$ an analytic function $\zeta : (-\delta, \delta) \rightarrow \mathbb{R}^K$ such that (i) $\zeta(0) = 0$; (ii) $(G, \sigma, \zeta(t)) \in h^i((G, \sigma) \times (0, 1))$ for all $t > 0$; and (iii) $o(\zeta) \leq \ell(\tau)$. Let $\eta = \sigma + \zeta$. Then η is a positive analytic function of order no more than $l(\tau)$ and against which σ is a best reply. To finish the proof, take l to be $\max_{\tau \in \mathcal{P}} l(\tau)$. □

References

Blume L, Brandenburger A, and Dekel E (1991a) Lexicographic Probabilities and Choice Under Uncertainty. *Econometrica* 59:61–79

- (1991b) Lexicographic Probabilities and Equilibrium Refinements. *Econometrica* 59:81–98
- Blume L, and Zame W (1994) The Algebraic Geometry of Perfect and Sequential Equilibrium. *Econometrica* 62:783–794
- Bochnak J, Coste M, and Roy M-F (1998) *Real Algebraic Geometry*. Springer-Verlag, Berlin
- Hammond P (1994) Elementary Non-Archimedean Representations of Probability for Decision Theory and Games. In *Patrick Suppes, Scientific Philosopher, Vol. 1*, ed. by Humphreys P, pp. 25–61. Kluwer Academic Press, Boston, MA
- Kohlberg E, and Reny P (1997) Independence on Relative Probability Spaces and Consistent Assessments in Game Trees. *Journal of Economic Theory* 75:280–313
- Kreps D, and Ramey G (1987) Structural Consistency, Consistency, and Sequential Rationality. *Econometrica* 55:1331–1348
- Kreps D, and Wilson R (1982) Sequential Equilibria. *Econometrica* 50:863–894
- McLennan A (1989) Consistent Conditional Systems in Noncooperative Game Theory. *International Journal of Game Theory* 18:141–174
- Myerson R (1978) Refinements of the Nash Equilibrium Concept. *International Journal of Game Theory* 7:73–80
- Selten R (1975) Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games. *International Journal of Game Theory* 4:25–55

Copyright of International Journal of Game Theory is the property of Springer Science & Business Media B.V.. The copyright in an individual article may be maintained by the author in certain cases. Content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.