A Nonstandard Characterization of Sequential Equilibrium, Perfect Equilibrium, and Proper Equilibrium*

Joseph Y. Halpern
Computer Science Department
Cornell University, U.S.A.
e-mail: halpern@cs.cornell.edu

August 11, 2008

Abstract

New characterizations of sequential equilibrium, perfect equilibrium, and proper equilibrium are provided that use nonstandard probability. It is shown that there exists a belief system μ such that $(\vec{\sigma}, \mu)$ is a sequential equilibrium in an extensive game with perfect recall iff there exist an infinitesimal ϵ and a completely mixed behavioral strategy profile σ' (so that σ'_i assigns positive, although possibly infinitesimal, probability to all actions at every information set) that differs only infinitesimally from $\vec{\sigma}$ such that at each information set I for player i, σ_i is an ϵ-best response to $\vec{\sigma}'_{-i}$ conditional on having reached I. Note that the characterization of sequential equilibrium does not involve belief systems. There is a similar characterization of perfect equilibrium; the only difference is that σ_i must be a best response to $\vec{\sigma}'_{-i}$ conditional on having reached I. Yet another variant is used to characterize proper equilibrium.

1 Introduction

Sequential equilibrium [Kreps and Wilson 1982] and *perfect equilibrium* [Selten 1975] are refinements of Nash equilibrium in extensive-form games defined in terms of "trembles" or mistakes. The definitions are subtle, involving sequences of strategies with smaller

*This work was supported in part by NSF under grants CTC-0208535, ITR-0325453, and IIS-0534064, and by AFOSR under grant FA9550-05-1-0055.

and smaller "trembles". In this paper I provide alternative characterizations of sequential equilibrium and perfect equilibrium using *nonstandard probability*. The definition replaces the sequence of trembles by a single (nonstandard) strategy profile and dispenses with belief systems altogether. I would argue that the new definition provides a better conceptual basis for the notion of sequential equilibrium and makes the relationship between sequential equilibrium, perfect equilibrium, and Nash equilibrium more transparent.

The new definition takes as its point of departure the observation that sequential equilibrium and perfect equilibrium have a simple definition for extensive-form games where the equilibrium is a completely mixed strategy profile, so that each strategy assigns positive probability to every action at every information set. A completely mixed strategy profile is a sequential equilibrium iff each player plays a best response *at all information sets*, that is, for each player i and each information set I for player i, player i plays a best response to the other players' strategies, conditional on having reached I. Moreover, a completely mixed strategy profile is (the strategy component of) a sequential equilibrium iff it is a perfect equilibrium.

All the subtlety in the definition of sequential equilibrium comes in dealing with information sets off the equilibrium path, that is, information sets that are reached with probability 0. To define "best response" at such an information set requires defining what the agent's beliefs are at that information set; this is exactly what is provided by the belief assessment. The consistency requirements imposed on belief assessments ensure that an agent's beliefs are in some sense "reasonable". At an information set that is reached with positive probability, there is an obvious way of defining "reasonable" using conditioning; indeed, an agent's beliefs at such an information set are completely determined by the strategies used. But what should count as a reasonable belief at an information set that the agent believes will not be reached in the first place? This concern is not new. Kreps and Wilson [1982, p. 876] themselves say: "We shall proceed here to develop the properties of sequential equilibrium as defined above; however, we do so with some doubts of our own concerning what 'ought' to be the definition of a consistent assessment that, with sequential rationality, will give the 'proper' definition of a sequential equilibrium." Osborne and Rubinstein [1994, p. 225] say "we do not find the consistency requirement to be natural, since it is stated in terms of limits; it appears to be a rather opaque technical assumption". They go on to quote Kreps [1990, p. 430], who says "[r]ather a lot of bodies are buried in this definition".

The new characterization avoids these difficulties to some extent by considering only completely mixed strategies. Intuitively, actions that are not best responses are given infinitesimal probability rather than 0 probability. As a first step towards the new characterization, consider the following characterization of sequential equilibrium, which is a variant of Kreps and Wilson's [1982] characterization of sequential equilibrium in terms of what they called *weak perfect equilibrium* (see their Proposition 6).

Proposition 1.1: *If Γ is an extensive-form game with perfect recall, then there exists a* *belief system μ such that the assessment $(\vec{\sigma}, \mu)$ is a sequential equilibrium in Γ iff there exists a sequence $\vec{\sigma}^n$ of completely mixed strategy profiles converging to $\vec{\sigma}$ and a sequence ϵ_n of nonnegative real numbers converging to 0 such that for each player i and each information set I for player i, σ_i^n is an ϵ_n-best response to $\vec{\sigma}_{-i}^n$, conditional on having reached I.*

Proposition 1.1 defines sequential equilibrium without needing assessments and is arguably quite natural. But it still defines a sequential equilibrium in terms of limits. Using nonstandard reals, we can replace the sequence $\vec{\sigma}^n$ of strategies in Proposition 1.1 by a single strategy and the sequence ϵ_n by a single infinitesimal to get a characterization that does not involve perturbations of either strategies or utility functions, and does not require assessments.

Theorem 1.2: *If Γ is an extensive-form game with perfect recall, then there exists a belief system μ such that the assessment $(\vec{\sigma}, \mu)$ is a sequential equilibrium in Γ iff there exist an infinitesimal ϵ and a nonstandard completely mixed strategy profile $\vec{\sigma}'$ that differs infinitesimally from $\vec{\sigma}$ such that, for each player i, and each information set I for player i, σ_i' is an ϵ-best response to $\vec{\sigma}'_{-i}$, conditional on having reached I.*

I define all the relevant notions in the statement of Theorem 1.2 in Section 2. I continue with a more informal discussion here, in the hope that the reader has enough of an intuitive sense of the relevant notions to follow.

One advantage of Theorem 1.2 is that it makes the connection between sequential equilibrium, Nash equilibrium, and perfect equilibrium more transparent. For Nash equilibrium, the following theorem, whose straightforward proof I leave to the reader, makes it clear that sequential equilibrium is a refinement of Nash equilibrium in the most obvious sense: we get a characterization of Nash equilibrium by dropping the requirement that σ_i be an ϵ-best response to $\vec{\sigma}_{-i}$ *at every information set*.

Theorem 1.3: *If Γ is an extensive-form game with perfect recall, then $\vec{\sigma}$ is a Nash equilibrium in Γ iff there exist an infinitesimal ϵ and a nonstandard completely mixed strategy profile $\vec{\sigma}'$ that differs infinitesimally from $\vec{\sigma}$ such that for each player i, σ_i' is an ϵ-best response to $\vec{\sigma}'_{-i}$.*

The connection to perfect equilibrium is equally straightforward. Note that in Theorem 1.2 I could have equally well required that σ_i be an ϵ-best response to $\vec{\sigma}'_{-i}$, since σ_i and σ_i' differ only infinitesimally. But if we focus on σ_i, we can ask for even more. We can require σ_i to be a *best response* to $\vec{\sigma}'_{-i}$, not just an ϵ-best response. As the following theorem shows, this strengthening of sequential equilibrium is perfect equilibrium.

Theorem 1.4: *If Γ is an extensive-form game with perfect recall, then $\vec{\sigma}$ is a perfect equilibrium in Γ iff there exists a nonstandard completely mixed strategy profile $\vec{\sigma}'$ that differs infinitesimally from $\vec{\sigma}$ such that, for each player i and each information set I for player i, σ_i is a best response to $\vec{\sigma}'_{-i}$, conditional on having reached I.*

This characterization of perfect equilibrium is not surprising; it just replaces the sequence of strategies in the standard definition of perfect equilibrium by a single strategy, just as in the transition from Proposition 1.1 and Theorem 1.2. However, it makes it clear that the only difference between sequential equilibrium and perfect equilibrium is whether σ_i is required to be a best response or an ϵ-best response to σ'_{-i}.

Proper equilibrium [Myerson 1978] further refines perfect equilibrium by, roughly speaking, requiring that bigger mistakes get smaller probability. This intuition can also be captured easily using nonstandard probability. Since proper equilibrium is typically defined for normal-form games, I give the characterization for normal form games.

Theorem 1.5: *If Γ is a normal-form game, then a mixed strategy profile $\vec{\sigma}$ is a proper equilibrium in Γ iff there exists a completely mixed nonstandard strategy profile $\vec{\sigma}'$ that differs infinitesimally from $\vec{\sigma}$ such that, for all players i, if s_1 and s_2 are two pure strategies for player i such that $\mathrm{EU}_i(s_1, \vec{\sigma}'_i) < \mathrm{EU}_i(s_2, \vec{\sigma}'_i)$ and $\mathrm{EU}_i(s_2, \vec{\sigma}'_i) - \mathrm{EU}_i(s_2, \vec{\sigma}'_i)$ is not infinitesimal, then $\sigma'_i(s_1)/\sigma'_1(s_2)$ is infinitesimal.*

It is easy to see that if $\vec{\sigma}$ is a proper equilibrium and $\sigma_i(s) > 0$, then s must be a best response to $\vec{\sigma}'_{-i}$. (Otherwise $\sigma'_i(s)$ would be infinitesimal, so $\sigma_i(s)$ would have to be 0, since $\sigma_i(s)$ and $\sigma'_i(s)$ differ only infinitesimally.) It follows that a proper equilibrium must be a perfect equilibrium.

Theorems 1.2, 1.4, and 1.5 serve to illustrate a well-known general phenomenon: sequences of real objects (strategies, numbers) can often by replaced by a single nonstandard object, leading to simpler statements of results. Of course, it may seem that, if the object is nonstandard, then there is no real gain in transparency or understanding. In practice, the sequence of "trembles" required to generate the off-equilibrium beliefs in a sequential equilibrium or the strategies converging to $\vec{\sigma}$ in a perfect or proper equilibrium is usually described in terms of a polynomial or rational function of ϵ, where ϵ goes to 0. Lemma 3.1 shows that this way of describing the trembles is almost without loss of generality; they can always be described using probabilities that are power series of ϵ (thus, we may need to go beyond rational functions to functions that involve sines, cosines, and exponentials). Moreover, it follows (Theorem 3.2) that the nonstandard probabilities can always be taken to be a power series in a single nonstandard ϵ. Thus, the nonstandard probabilities involved are not so mysterious.[1]

Theorem 1.2 also suggests an alternative interpretation of sequential equilibrium. It shows that (the strategy component of) a sequential equilibrium is the standard part of an ϵ-sequential equilibrium for an infinitesimal ϵ. (For every nonstandard number r, there is a unique closest standard real number r', which differs from r by an infinitesimal; r' is called the *standard part* of r. We can then define the standard part of a nonstandard behavioral strategy profile in the obvious way.) Under this interpretation, the focus is

[1]The suggestion to prove Lemma 3.1 and Theorem 3.2, as well as sketches of their proof, were provided by an anonymous reviewer.

on the ϵ-sequential equilibrium $\vec{\sigma}'$, not $\vec{\sigma}$; $\vec{\sigma}$ is just the closest standard approximation to the "true" equilibrium.

An interpretation of Theorems 1.2 and 1.4 that is somewhat closer to the standard interpretation is that they show that in a sequential equilibrium (resp., perfect equilibrium) $\vec{\sigma}$, each player i can be viewed as making an ϵ-best response (resp., best response) at each information set, under the belief that the other players are playing $\vec{\sigma}'_{-i}$ (so that player i's beliefs off the equilibrium path are determined by $\vec{\sigma}'_{-i}$). This belief is reasonable: even repeated observations of the game will not contradict it (the distribution induced by $\vec{\sigma}$ is essentially indistinguishable from that induced by $\vec{\sigma}'$). However, this interpretation has an obvious difficulty: why should we be interested in the particular beliefs induced by $\vec{\sigma}'$? Although they are reasonable, why not consider the beliefs induced by other strategies $\vec{\sigma}''$ that are also infinitesimally close to $\vec{\sigma}$? Put another way, why is it good enough to know that $\vec{\sigma}$ is an (ϵ-)best response to *some* beliefs? In the more standard definition of sequential equilibrium, this question can be reformulated as asking why we should care that σ_i is a best response to the particular beliefs given by the belief assessment. The advantage of the first interpretation is that it avoids this issue altogether.

Related work: As I said, Proposition 1.1 is much in the spirit of Kreps and Wilson's characterization of sequential equilibrium in terms of weak perfect equilibrium. Rather than comparing Proposition 1.1 to Kreps and Wilson's results, I compare it to a reformulation due to Blume and Zame [1994]. In their Proposition A, Blume and Zame show that $(\vec{\sigma}, \mu)$ is a sequential equilibrium in a game Γ with utility functions $\vec{u}$ iff there exists a sequence $\vec{\sigma}^n$ of completely mixed strategies converging to $\vec{\sigma}$ and a sequence $\vec{u}^n$ of utility functions converging to $\vec{u}$ such that for each player i and each information set I for player i, σ_i^n is a best response to $\vec{\sigma}^n_{-i}$, conditional on having reached I, with respect to the utility function $\vec{u}^n$. Clearly if, for all n, σ_i^n is a best response to $\vec{\sigma}^n_{-i}$, conditional on having reached I, with respect to the utility function $\vec{u}^n$, then we can find a sequence ϵ^n converging to 0 such that σ_i^n is an ϵ^n-best response to $\vec{\sigma}^n_{-i}$ conditional on having reached I with respect to $\vec{u}$. Conversely, if the condition stated in Proposition 1.1 holds, it can be shown that Blume and Zame's condition holds as well, although it is actually easier to show directly that there is a sequential equilibrium. Since the characterization of Theorem 1.2 does not need to use sequences of trembles and does not require modifying the game, it is arguably more natural, although "naturalness" is clearly in the eye of the beholder here.

The results of this paper are far from the first use of nonstandard analysis in game theory. Perhaps closest to the results of this paper, Blume, Brandenburger, and Dekel [1991] characterize perfect and proper equilibrium using *lexicographic probability systems* (LPS's). There are deep connection between LPS's and nonstandard probability. Hammond [1994] proves that they are equivalent in finite spaces; his results are generalized in [Halpern 2001]. Using these connections, it is not hard to prove Theorems 1.4 and 1.5 from the characterizations given by Blume, Brandenburger, and Dekel, and vice versa. More generally, Hammond [1994] argued for thinking of notions such as perfect and proper equilibrium in terms of nonstandard probabilities. Rajan [1998] characterizes a number of solution concepts, including perfect equilibrium and proper equilibrium (but not sequential equilibrium), in terms of agents whose beliefs are given by nonstandard probability distributions. Nonstandard utilities and lexicographic utilities have also been considered by a number of authors (see, e.g., [Fishburn 1972; Richter 1971; Skala 1974] and the references therein).

Outline: The rest of this paper is organized as follows: In Section 2 I give all the relevant formal definitions (and include a review of the definition of sequential equilibrium). Theorems 1.2 and 1.4 are proved in Section 3. These proofs depend on a deep theorem of first-order logic, the *compactness theorem* [Enderton 1972], which is explained in Section 2. The proof of Theorem 1.5 is so similar to the other two that it is omitted; the proof of Theorem 1.1 is similar in spirit to, but simpler than, that of Theorem 1.2, so is also omitted.

2 Definitions

To make the paper self-contained, I briefly review the definitions of sequential equilibrium, perfect equilibrium, and proper equilibrium, and then discuss relevant concepts from logic. I assume that the reader is familiar with the standard definition of games in extensive form and normal form. All of the notions I consider are standard in the game theory and logic literatures.

2.1 Sequential equilibrium

Sequential equilibrium is defined with respect to an *assessment*, a pair $(\vec{\sigma}, \mu)$ where $\vec{\sigma}$ is a strategy profile and μ is a *belief system*, i.e., a function that determines for every information set I a probability μ_I over the histories in I. Intuitively, if I is an information set for player i, then μ_I is i's subjective assessment of the relative likelihood of the histories in I. Roughly speaking, an assessment is a sequential equilibrium if (a) at every information set where a player moves he chooses a best response given the beliefs he has about the histories in that information set and the strategies of other players, and (b) his beliefs are consistent with the strategy profile being played. Consistency at an information set I is easy to define if I is reached with positive probability by $\vec{\sigma}$; in that case, it is just the result of conditioning on I. The definition is a little more subtle if I is reached with probability 0.

Given a strategy profile $\vec{\sigma}$, let $\Pr_{\vec{\sigma}}$ be the probability distribution on complete histories of the game induced by $\vec{\sigma}$. That is, if h is a complete history, then $\Pr_{\vec{\sigma}}(h)$ is the product of the probability of each of the moves made in h. If we identify a partial history with the set of complete histories that extend it, this statement holds for partial histories as well.

If I is an information set, we take $\Pr_{\vec{\sigma}}(I)$ to be the probability of the set of complete histories extending some partial history in I. The conditional probability $\Pr_{\vec{\sigma}}(h \mid I)$ is defined in the standard way.

Formally, an assessment $(\vec{\sigma}, \mu)$ is a *sequential equilibrium* of an extensive-form game Γ if it satisfies the following properties:

- *Sequential rationality.* For every information set I of player i and every behavioral strategy τ for player i, σ_i is a best response to $\vec{\sigma}_{-i}$ at I given belief system μ; that is,
$$\mathrm{EU}_i((\vec{\sigma}, \mu) \mid I) \geq \mathrm{EU}_i(((\tau, \vec{\sigma}_{-i}), \mu) \mid I),$$
where $\mathrm{EU}_i((\vec{\sigma}, \mu) \mid I) = \sum_{h \in I} \sum_{z \in Z} \mu_I(h) \Pr_{\vec{\sigma}}(z \mid h) u_i(z)$ and Z is the set of terminal histories of the game Γ.

- *Consistency between belief system and strategy profile.* There exists a sequence of assessments $(\vec{\sigma}^n, \mu^n)$, $n = 1, 2, 3, \ldots$, such that, for all n, $\vec{\sigma}^n$ is a profile of completely mixed strategies, $\vec{\sigma}^n \to \vec{\sigma}$, and $\mu_I(h) = \lim_{n \to \infty} \Pr_{\vec{\sigma}^n}(h \mid I)$. (Since each strategy in $\vec{\sigma}^n$ is completely mixed, $\Pr_{\vec{\sigma}^n}(I) > 0$, so the conditional probability is well defined.)

2.2 Perfect equilibrium and proper equilibrium

To define perfect equilibrium, I first need to define the notion of a best response at an information set.

Definition 2.1: If $\epsilon \geq 0$ and I is an information set for player i, σ_i is an *ϵ-best response to $\vec{\sigma}_{-i}$ for i conditional on having reached I* if for every strategy τ for player i that agrees with σ except possibly at I and information sets preceded by I we have $\mathrm{EU}_i(\vec{\sigma}) \geq \mathrm{EU}_i(\tau, \vec{\sigma}_{-i}) - \epsilon$. A *best response* is a 0-best response. ∎

Definition 2.2: $\vec{\sigma}$ is a *perfect equilibrium* in an extensive game Γ iff there exists a sequence $\vec{\sigma}^n$ of completely mixed behavior strategies such that $\vec{\sigma}^n \to \vec{\sigma}$ and, for all n and each information set I of player i, σ_i is a best response to $\vec{\sigma}^n_{-i}$ conditional on having reached I.[2] ∎

Proper equilibrium is typically defined in normal-form games. As I said earlier, proper equilibrium tries to capture the intuition that worse strategies get smaller probability.

[2]It is more standard to define perfect equilibrium as above in normal-form games, and then define a perfect equilibrium in an extensive-form game Γ to be a perfect equilibrium of the normal-form game Γ' corresponding to the *agent-normal* form of Γ, where the agent-normal form of Γ has one agent for each information set in Γ; the agents corresponding to information sets of player i in Γ make decisions independently but have the same utility as player i. It is not hard to show (using arguments similar to those of Selten [1975, Lemma 6]) that, in games with perfect recall, the standard definition is equivalent to the one given here.

Definition 2.3: $\vec{\sigma}$ is a *proper equilibrium* in a normal-form game if there exists a sequence of completely mixed strategy $\vec{\sigma}^n$ and a sequence of positive real numbers $\epsilon_1, \epsilon_2, \ldots$ such that $\vec{\sigma}^n \to \vec{\sigma}$, $\epsilon_n \to 0$, and, for all players i and all pure strategies s_1 and s_2 for player i, if $\mathrm{EU}_i(s_1, \vec{\sigma}^k_{-i}) < \mathrm{EU}_i(s_2, \vec{\sigma}^k_{-i})$, then $\sigma_i(s_1) < \epsilon_k \sigma_i(s_2)$. ∎

2.3 Nonstandard Probabilities

An *elementary extension F of the reals* is an ordered field that includes the real numbers, at least one *infinitesimal* (i.e., a number ϵ that is less than every positive real number, but still greater than 0), and is *elementarily equivalent to the real numbers*. The fact that F and $\mathbb{R}$ are elementarily equivalent means that every formula that can be expressed in first-order logic and uses the function symbols $+$ and $\times$ and a constant symbol $\mathbf{r}$ for each real number r is true in F iff it is true in $\mathbb{R}$ (where, in $\mathbb{R}$, $+$ is interpreted as addition, $\times$ is interpreted as multiplication, and the constant symbol $\mathbf{r}$ is interpreted as the real number r).

The existence of such elementary extensions of the reals follows from the *compactness theorem* of first-order logic. The compactness theorem says that if every finite subset of an infinite collection of formulas has a model, then the whole infinite set has a model. To see how the existence of elementary extensions of the reals follows from the compactness theorem, add an extra constant symbol $\mathbf{c}$ to the language of the reals, and consider the (uncountable) set Φ of formulas consisting of

(a) all first-order formulas in the language true of the reals (this includes, for example, a formula such as $\forall x \forall y (x + y = y + x)$, which says that addition is commutative, as well as formulas such as $\mathbf{2} + \mathbf{3} = \mathbf{5}$ and $\sqrt{\mathbf{2}} \times \sqrt{\mathbf{3}} = \sqrt{\mathbf{6}}$);

(b) the countable collection of formulas $\mathbf{c} < 1/\mathbf{n}$, for each positive natural n;

(c) the formula $\mathbf{c} > 0$.

Clearly each finite subset Φ' of Φ has a model; namely, the reals; there is always an interpretation for $\mathbf{c}$ as a real number if only finitely many statements in (b) are considered. Compactness says that Φ must have a model, say F. In F, $\mathbf{c}$ must be interpreted as an infinitesimal, so F must be an elementary extension of $\mathbb{R}$.

If F is an elementary extension of $\mathbb{R}$, I refer to the elements of F that are in $\mathbb{R}$ as standard reals. It is not hard to show that for all $b \in F$ such that $-r < b < r$ for some standard real $r > 0$, there is a unique closest real number a such that $|a - b|$ is an infinitesimal. (In fact, a is the inf of the set of real numbers that are at least as large as b.) Let $st(b)$ denote the closest standard real to b; $st(b)$ is sometimes read "the standard part of b".

For the purposes of this paper, I take a nonstandard probability to be one whose range lies in an elementary extension of the reals. Two nonstandard probabilities Pr and

Pr$'$ are *infinitesimally close* if, for each event E, $|\Pr(E) - \Pr'(E)|$ is an infinitesimal. Given a nonstandard probability measure Pr, there is a unique standard probability measure Pr$'$ that is infinitesimally close to it, obtained by taking $\Pr'(E) = st(\Pr(E))$. A behavior strategy for player i just defines a probability distribution over actions at each information set of player I. A *nonstandard* behavior strategy is one that uses a nonstandard probability distribution; a *standard behavior strategy* is one that uses a standard probability distribution. Two behavior strategies σ and σ' for player i are infinitesimally close if, for every information set I of player i, the probability distributions defined by σ and σ' are infinitesimally close.

In Section 3.4, I show that in all the main theorems we can use one particular, quite natural, elementary extension of the reals, denoted $I\!R^*(\epsilon)$. To describe $I\!R^*(\epsilon)$, first consider $I\!R(\epsilon)$, the smallest non-Archimedean field strictly containing the reals. $I\!R(\epsilon)$ consists of all rational expressions $f(\epsilon)/g(\epsilon)$ in ϵ, where f and g are polynomials and $g \neq 0$. (The construction of $I\!R(\epsilon)$ apparently goes back to Robinson [1973]; $I\!R(\epsilon)$ is used extensively by Hammond [1994].) While $I\!R(\epsilon)$ is an ordered non-Archimedean field, it is not an elementary extension of $I\!R$. For example, ϵ does not have a cube root in $I\!R(\epsilon)$, although it is a theorem of $I\!R$ that every number has a cube root. As a first step to obtaining an elementary extension of $I\!R$, let $I\!R^+(\epsilon)$ consist of all formal power series in ϵ of the form $\sum_{i \geq n} r_i \epsilon^i$, where n is an integer and $r_i \in I\!R$. $I\!R^+(\epsilon)$ is also easily seen to be an ordered non-Archimedean field, where inverses are computed by formal division of power series. (Note that we need to allow the summation to start at a negative number since, for example, the inverse of ϵ^3 is ϵ^{-3}.) Although $I\!R^+(\epsilon)$ includes $I\!R(\epsilon)$ (every rational function can be expressed as a power series by formal division), it is still not an elementary extension of $I\!R$; again ϵ has no cube root. In order to extend $I\!R^+(\epsilon)$ to an elementary extension of $I\!R$, we need a field that includes, for example $\epsilon^{1/3}$, so that ϵ has a cube root. This suggests that, in order to construct an elementary extension of $I\!R(\epsilon)$, it is necessary to allow rational exponents in power series. Let $I\!R^*(\epsilon)$ consist of all series of the form $\sum_{i \geq n} r_i \epsilon^{i/k}$, where n is an integer and and k is a positive natural number. $I\!R^*(\epsilon)$ is known as the field of *Puiseux series in* ϵ, and is an elementary extension of $I\!R$ [Basu, Pollack, and Roy 2003, Chapter 2].

3 Proofs of Theorems

3.1 Proof of Proposition 1.1

Suppose that the assessment $(\vec{\sigma}, \mu)$ is a sequential equilibrium. Then there exists a sequence of assessments $(\vec{\sigma}^n, \mu^n)$ converging to $(\vec{\sigma}, \mu)$. It is easy to see that there must be a sequence ϵ_n of nonnegative real numbers converging to 0 such that $\vec{\sigma}^n$ is an ϵ_n-sequential equilibrium.

For the converse, if $\vec{\sigma}^n$ is a sequence of completely mixed strategy profiles converging to $\vec{\sigma}$, then let μ^n be the belief system determined by $\vec{\sigma}^n$. By compactness, there must be a convergent subsequence of μ^n. Let μ be the belief system to which this convergent subsequence converges. It easily follows from the fact that $\vec{\sigma}^n$ is an ϵ_n-sequential equilibrium for some sequence of nonnegative reals ϵ_n converging to 0 that $(\vec{\sigma}, \mu)$ is a sequential equilibrium.

3.2 Proof of Theorem 1.2

First suppose that $\vec{\sigma}'$ satisfies the conditions of the theorem. Define a belief system μ by taking $\mu_I(h) = st(\Pr_{\vec{\sigma}}(h \mid I))$; that is, the standard part of the conditional probability of h given I, according to $\Pr_{\vec{\sigma}'}$. Since $\vec{\sigma}'$ consists of completely mixed strategies, this conditional probability is well defined. We want to show that $(\vec{\sigma}, \mu)$ is a sequential equilibrium. Because there exists an infinitesimal ϵ such that σ_i is an ϵ-best response to $\vec{\sigma}'_{-i}$ conditional on having reached I, and $\vec{\sigma}'$ differs infinitesimally from $\vec{\sigma}$, it easily follows that, for all standard strategies τ for player i, we have $\mathrm{EU}_i((\vec{\sigma}, \mu) \mid I) \geq \mathrm{EU}_i(((\tau, \vec{\sigma}_{-i}), \mu) \mid I)$; thus, sequential rationality holds. To show that we have consistency between the belief system and strategy profile, we use the fact that F, the field which is the range of the nonstandard probabilities used in $\vec{\sigma}'$, is an elementary extension of $I\!R$. The idea is, for each n, to write a first-order formula φ_n that says that a strategy profile $\vec{\sigma}^n$ with the right properties exists. Since φ_n will be easily seen to be true in F, it must also be true in $I\!R$. In more detail, we proceed as follows.

For each information set I of player i, if h is in I, let $q_h = \Pr_{\vec{\sigma}'}(h|I)$; let $r_h = \mu_I(h) = st(q_h)$. In addition, for each action a that i can perform in information set I, let $q_{I,a} = \sigma'_i(I)(a)$; that is, $q_{I,a}$ is the probability that i performs action a in information set I according to σ'_i. Let $r_{I,a} = \sigma_i(I)(a) = st(q_{I,a})$. Suppose that there are M histories in the game Γ, K information sets in Γ, and for each information set I_j in Γ, there are N_j actions that can be performed at I_j. Now, for each n, let φ_n be a formula $\exists x_{h_1} \ldots \exists x_{h_M} \exists x_{I_1,a_{11}} \ldots \exists x_{I_1,a_{1N_1}} \ldots \exists x_{I_K,a_{K1}} \ldots \exists x_{I_K,a_{KN_K}} \varphi'_n$, where φ'_n is described below. Intuitively, φ'_n says that x_h, the conditional probability of h given I (the information set containing h) is within $1/n$ of q_h, and that $x_{I_j,a_{jk}}$, the probability of performing action a_{jk} in information set I_j, is within $1/n$ of $q_{I_j,a_{jk}}$. Finally, for each history h in information set I, φ'_n relates x_h to probability of reaching h conditional on reaching I. Thus, φ'_n is the conjunction of the following formulas:

(a) formulas of the form $|x_h - \mathbf{r}_h| < 1/n$ for all histories h;

(b) formulas of the form $(|x_{I,a} - \mathbf{r}_{I,a}| < 1/n) \wedge (x_{I,a} > 0)$ for each information I and action a that can be performed at I;

(c) for each information set I_j, a formulas $\sum_{k=1}^{N_j} x_{I_j,a_{jk}} = 1$;

(d) for each history h in an information set I, a formula saying that x_h is the product of the probabilities $x_{I',a}$ over the actions a taken to reach h, divided by the sum of the corresponding products taken over all histories $h' \in I$.

The formula φ_n is clearly satisfiable in F by taking x_h to be q_h and $x_{I,a}$ to be $q_{I,a}$. Thus, φ_n is also satisfiable in $I\!R$. The values of $x_{I,a}$ that satisfy the formula in $I\!R$ determine the strategy profile $\vec{\sigma}^n$; if I is an information set of player I, then $\sigma_i^n(I)(a) = x_{I,a}$. The formulas in (b) and (c) guarantee that σ_i^n is a completely mixed strategy that is within $1/n$ of σ_i. Thus $\vec{\sigma}^n \to \vec{\sigma}$. Taking $\mu_I^n(h) = x_h$, the formulas in clause (d) guarantee that μ_I^n is the (unique) belief assessment corresponding to the strategy profile $\vec{\sigma}^n$ (it is unique because $\vec{\sigma}^n$ consists of completely mixed strategies). The formulas in clause (a) guarantee that μ^n is within $1/n$ of μ, so $\mu^n \to \mu$.

For the converse, suppose that $(\vec{\sigma}, \mu)$ is a sequential equilibrium. Thus, there exists a sequence of assessments $(\vec{\sigma}^n, \mu^n) \to (\vec{\sigma}, \mu)$ such that, for each n, $\vec{\sigma}^n$ is a profile of completely mixed strategies. Since, at each information set I for player i, σ_i is a best response to $\vec{\sigma}_{-i}$ conditional on having reached I, given belief system μ, it must be the case that, for all ϵ, there exists n such that σ_i^n is an ϵ-best response to $\vec{\sigma}_{-i}^n$ at I, given belief system μ^n. But since $\vec{\sigma}^n$ consists of completely mixed strategies, so that μ^n is determined by $\vec{\sigma}^n$, it is easy to see that σ_i^n is in fact an ϵ-best response to σ_{-i}^n conditional on having reached I, in the sense of Definition 2.1.

With this observation, we can now use the compactness theorem to show that an appropriate $\vec{\sigma}'$ exists. The proof is similar in spirit to that of the first half. We consider the language of the reals with additional constants $\mathbf{c}_h$ for each history h, a constant $\mathbf{c}_{I,a}$ for each (information set, action) pair, and a constant $\mathbf{d}$. Consider the following collection Φ of formulas:

(a) all first-order formulas in the language true of the reals;

(b) for each information I and action a that can be performed at I, the formula $\mathbf{c}_{I,a} > 0$ and, for each natural number $n > 0$, the formula $|\mathbf{c}_{I,a} - \mathbf{r}_{I,a}| < 1/\mathbf{n}$;

(c) for each information set I_j, a formula $\sum_{k=1}^{N_j} \mathbf{c}_{I_j,a_{jk}} = 1$;

(d) the formula $\mathbf{d} > 0$ and, for each natural number the formula $\mathbf{d} < 1/\mathbf{n}$;

(e) for each player i, a formula saying that the strategy σ_i is a $\mathbf{d}$-best response to the strategy profile $\vec{\sigma}_{-i}$ defined by the probabilities $\mathbf{c}_{I,a}$. The statement that σ_i is a $\mathbf{d}$-best response is easily definable using a universally-quantified formula with linear inequalities.

Any finite subset Φ' of Φ is satisfiable by the reals. Since F contains only finitely many formulas of the form $\mathbf{d} < 1/n$, there is some real number ϵ that satisfies these formulas. We can then choose $\mathbf{c}_{I,a}$ to be the probabilities determined by $\vec{\sigma}^n$ for n sufficiently large so that, for each information set I for player i, σ_i^n is an ϵ-best response to $\vec{\sigma}_{-i}^n$. That is, for the appropriate choice of n, if I is an information set for player i, then $(c)_{I,a}$ is interpreted as $\sigma_i^n(I)(c)$.

Since every finite subset of Φ is satisfiable, by compactness, there is a model F for Φ. The formulas in (a) force F to be an elementary extension of $I\!R$; the formulas in

(c) ensure that the interpretation of the constants $\mathbf{c}_{I,a}$ determines a strategy profile $\vec{\sigma}'$; the formulas in (b) guarantee that $\vec{\sigma}$ and $\vec{\sigma}'$ differ only infinitesimally; the formulas in (d) guarantee that the interpretation ϵ of $\mathbf{d}$ is an infinitesimal; and the formulas in (e) guarantee that σ_i is an ϵ-best response to $\vec{\sigma}'$. This proves the desired result.

3.3 Proof of Theorem 1.4

The structure of the proof of Theorem 1.4 is almost identical to that of Theorem 1.2, so here I just highlight the differences.

Suppose that there exists a completely mixed nonstandard strategy profile $\vec{\sigma}'$ that differs infinitesimally from $\vec{\sigma}$ such that at each information set I for player i, σ_i is a best response to $\vec{\sigma}'_{-i}$, conditional on having reached I. For each n, we can define a formula ψ_n that can be used to define the strategy profile $\vec{\sigma}^n$; ψ_n has same form as the formula φ_n in the proof of Theorem 1.2. We just need to add another finite collection of conjunctions to the formula φ_n' saying that, at each information set I of player i, σ_i is a best response to the strategy profile defined by the values $x_{I,a}$. Again, the formula φ_n is satisfied in F, the field which is the range of the nonstandard probabilities used in $\vec{\sigma}'$ by interpreting $x_{I,a}$ as $q_{I,a}$ and x_h as q_h. Thus, φ_n is satisfied in $I\!R$. The interpretation of the variables $x_{I,a}$ in $I\!R$ determines the strategy $\vec{\sigma}^n$.

For the converse, we again take an infinite collection of formulas and apply compactness. The formulas are identical to those used in the proof of Theorem 1.2, except that we do not need the constant $\mathbf{d}$ and the constraints on $\mathbf{d}$ described by the formulas in (d), and we modify the formulas in (e) to say that that σ_i is a best response to the strategy profile defined by the probabilities $\mathbf{c}_{I,a}$, rather than just a $\mathbf{d}$-best response.

3.4 A Power Series Representation of Nonstandard Strategy Profiles

As I said in the introduction, Theorem 1.2, 1.4, and 1.5 do not give any hint as to how to think of the nonstandard strategy profiles $\vec{\sigma}'$ that differ infinitesimally from $\vec{\sigma}$. In this section, I show that they can be thought of in a way that is close in spirit to the way trembles are typically presented in the literature; we can take all the probabilities to lie in $I\!R^*(\epsilon)$.[3]

As a first step, we need the following lemma. Note that a standard strategy in a (finite) extensive-form game can be identified with a vector in $I\!R^N$ for an appropriate choice of N, since the strategy is just characterized by a tuple of numbers describing the probabilities of moves at each information set. Of course, the same is true for a normal-form game as well. In the following lemma, I assume that N is chosen appropriately. Recall that an analytic function is one that is locally given by a convergent power series.

[3]I think Hari Govindan both for suggesting that I prove this result and for providing a sketch of the proof.

Lemma 3.1: *Suppose that $\vec{\sigma}$ is a (standard) strategy profile.*

(a) $\vec{\sigma}$ is a sequential equilibrium of an extensive-form Γ with perfect recall iff there exists analytic functions $\bar{\sigma} : [0,1] \rightarrow I\!R^N$ and $f : [0,1] \rightarrow [0,1]$ such that $\bar{\sigma}(0) = \vec{\sigma}$, $f(0) = 0$, and, for all $\epsilon > 0$, $\bar{\sigma}(\epsilon)$ is a completely mixed strategy profile such that, for each player i and each information set I for player i, $\bar{\sigma}_i(\epsilon)$ is an $f(\epsilon)$-best response to $\bar{\sigma}_{-i}(\epsilon)$, conditional on having reached I.

(b) $\vec{\sigma}$ is a perfect equilibrium of an extensive-form game Γ with perfect recall iff there exists an analytic functions $\bar{\sigma} : [0,1] \rightarrow I\!R^N$ such that $\bar{\sigma}(0) = \vec{\sigma}$ and, for all $\epsilon > 0$, $\bar{\sigma}(\epsilon)$ is a completely mixed strategy profile such that, for each player i and each information set I of player i, σ_i is a best response to $\bar{\sigma}_{-i}(\epsilon)$ conditional on having reached I.

(c) $\vec{\sigma}$ is a proper equilibrium of a normal-form game Γ iff there exist analytic functions $\bar{\sigma} : [0,1] \rightarrow I\!R^N$ and $f : [0,1] \rightarrow [0,1]$ such that $\bar{\sigma}(0) = \vec{\sigma}$, $f(0) = 0$, and, for all $\epsilon > 0$, $\bar{\sigma}(\epsilon)$ is a completely mixed strategy profile such that each player i and all pure strategies s_1 and s_2 of player i, if $\mathrm{EU}_i(s_1, \bar{\sigma}_{-i}(\epsilon) < \mathrm{EU}_i(s_2, \bar{\sigma}_{-i}(\epsilon)$, then $\bar{\sigma}_i(s_1) < f(\epsilon)\bar{\sigma}_i(s_2)$.

Proof: Sufficiency in all three cases is clear (for part (a), this depends on Proposition 1.1). For necessity, I first prove part (b). Let A be the set of completely mixed strategy profiles $\vec{\tau}$ such that, for each player i and each information set I of player i, σ_i is a best response to $\vec{\tau}_{-i}$ conditional on having reached I. A is easily seen to be a *semialgebraic* set (that is, a set defined by a finite sequence of polynomial equations and inequalities, or the union of such sets); moreover, it is immediate from the definition of perfect equilibrium that $\vec{\sigma}$ is in the closure of A. The result now follows immediately from the Nash Curve Selection Lemma [Bochnak, Coste, and Roy 1998, Proposition 8.1.13].

The modifications of this results to prove parts (a) and (c) are straightforward. For (a), let A consist of pairs $(\vec{\tau}, \epsilon)$ such that, for each player i and each information set I for player i, τ_i is an ϵ-best response to $\vec{\tau}_{-i}$, conditional on having reached I; for (c), let A consist of pairs $(\vec{\tau}, \epsilon)$ such that if $\mathrm{EU}_i(s_1, \vec{\tau}_{-i}(\epsilon) < \mathrm{EU}_i(s_2, \vec{\tau}_{-i}(\epsilon)$, then $\tau_i(s_1) < \epsilon\tau_i(s_2)$. In both cases, $(\vec{\sigma}, 0)$ is in the closure of A (for part (a), this follows from Proposition 1.1). Again, the result follows from the Nash Curve Selection Lemma. ∎

Theorem 3.2: *In Theorems 1.2, 1.4, and 1.5, the extension field F can be taken to be $I\!R^*(\epsilon)$, and all the probabilities can be taken to be analytic functions of ϵ (i.e., power series of the form $\sum_{i\geq 0} r_i\epsilon^i$).*

Proof: The fact that the probabilities can be taken to be analytic functions of ϵ in Theorems 1.2, 1.4, and 1.5 is immediate from Lemma 3.1. We can simply taken $\vec{\sigma}'$ in each of these theorems to be $\bar{\sigma}(\epsilon)$. Since $\bar{\sigma}(0) = \vec{\sigma}$, it follows that each component of

$\vec{\sigma}' - \vec{\sigma}$ is a power series whose leading coefficient is 0, and hence is infinitesimal. For (a) and (c), since $f(0) = 0$, $f(\epsilon)$ is an infinitesimal.

The fact that $\vec{\sigma}$ is a sequential (resp., perfect, proper) equilibrium if the conditions in Theorem 1.2 (resp., Theorem 1.4, Theorem 1.5) hold with for a nonstandard mixed strategy $\vec{\sigma}$ whose probabilities are all analytic functions of ϵ is immediate from the proofs of these theorems (where it is shown that these results hold as long as the probabilities come from any elemenatry extension of $I\!R$). ∎

Acknowledgments

Thanks to Adam Brandenburger, Larry Blume, Ehud Kalai, and Martin Osborne for useful discussions. Thanks also to Adam for suggesting that I state Proposition 1.1, to make the result more palatable to those unused to infinitesimals, and to Larry for pointing out the connections between this result and Proposition A in [Blume and Zame 1994]. Finally, thanks to Hari Govindan for suggesting that I prove Theorem 3.2 and providing a sketch proof.

References

Basu, S., R. Pollack, and M.-F. Roy (2003). *Algorithms in Real Algebraic Geometry*. Algorithms and Computation in Mathematics, Volume 10. Berlin/New York: Springer-Verlag.

Blume, L., A. Brandenburger, and E. Dekel (1991). Lexicographic probabilities and equilibrium refinements. *Econometrica 59*(1), 81–98.

Blume, L. E. and W. R. Zame (1994). The algebraic geometry of perfect and sequential equilibrium. *Econometrica 62*, 783–794.

Bochnak, J., M. Coste, and M.-F. Roy (1998). *Real Algebraic Geometry*. Berlin/New York: Springer-Verlag.

Enderton, H. B. (1972). *A Mathematical Introduction to Logic*. New York: Academic Press.

Fishburn, P. C. (1972). On the foundations of game theory: the case of non-Archimedean utilities. *International Journal of Game Theory 1*, 65–71.

Halpern, J. Y. (2001). Lexicographic probability, conditional probability, and nonstandard probability. In *Theoretical Aspects of Rationality and Knowledge: Proc. Eighth Conference (TARK 2001)*, pp. 17–30. San Francisco: Morgan Kaufmann. A full version of the paper can be found at http://arxiv.org/abs/cs/0306106.

Hammond, P. J. (1994). Elementary non-Archimedean representations of probability for decision theory and games. In P. Humphreys (Ed.), *Patrick Suppes: Scientific Philosopher; Volume 1*, pp. 25–49. Dordrecht, Netherlands: Kluwer.

Kreps, D. M. (1990). *Game Theory and Economic Modeling.* Oxford, UK: Oxford University Press.

Kreps, D. M. and R. B. Wilson (1982). Sequential equilibria. *Econometrica 50*, 863–894.

Myerson, R. (1978). Refinements of the Nash equilibrium concept. *International Journal of Game Theory 7*, 73–80.

Osborne, M. J. and A. Rubinstein (1994). *A Course in Game Theory.* Cambridge, Mass.: MIT Press.

Rajan, U. (1998). Trembles in the Bayesian foundation of solution concepts. *Journal of Economic Theory 82*, 248–266.

Richter, M. K. (1971). Rational choice. In J. Chipman, L. Hurwicz, M. K. Richter, and H. F. Sonnenschein (Eds.), *Prefrence, Utility, and Demand*, pp. 29–58. New York: Harcourt, Brace, and Jovanovich.

Robinson, A. (1973). Function theory on some nonarchimedean fields. *American Mathematical Monthly: Papers on the Foundations of Mathematics 80*, S87–S109.

Selten, R. (1975). Reexamination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory 4*, 25–55.

Skala, H. J. (1974). Nonstandard utilities and the foundations of game theory. *International Journal of Game Theory 3*(2), 67–81.