

Characterizing Solution Concepts in Terms of Common Knowledge of Rationality

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

Yoram Moses†
Department of Electrical Engineering
Technion—Israel Institute of Technology
32000 Haifa, Israel
email: moses@ee.technion.ac.il

March 5, 2016

Abstract

Characterizations of Nash equilibrium, correlated equilibrium, and rationalizability in terms of common knowledge of rationality are well known [Aumann 1987; Brandenburger and Dekel 1987]. Analogous characterizations of sequential equilibrium, (trembling hand) perfect equilibrium, and quasi-perfect equilibrium in n -player games are obtained here, using results of Halpern [2009, 2013].

*Supported in part by NSF under grants CTC-0208535, ITR-0325453, and IIS-0534064, by ONR under grant N00014-02-1-0455, by the DoD Multidisciplinary University Research Initiative (MURI) program administered by the ONR under grants N00014-01-1-0795 and N00014-04-1-0725, and by AFOSR under grants F49620-02-1-0101 and FA9550-05-1-0055.

†The Israel Pollak academic chair at the Technion; work supported in part by Israel Science Foundation under grant 1520/11.

1 Introduction

Arguably, the major goal of epistemic game theory is to characterize solution concepts epistemically. Characterizations of the solution concepts that are most commonly used in strategic-form games, namely, Nash equilibrium, correlated equilibrium, and rationalizability, in terms of common knowledge of rationality are well known [Aumann 1987; Brandenburger and Dekel 1987]. We show how to get analogous characterizations of sequential equilibrium [Kreps and Wilson 1982], (trembling hand) perfect equilibrium [Selten 1975], and quasi-perfect equilibrium [van Damme 1984] for arbitrary n -player games, using results of Halpern [2009, 2013].

To put our results in context, we start by reviewing the characterizations of Nash equilibrium, correlated equilibrium, and rationalizability in Section 2. In Section 3, we recall Halpern’s characterizations of sequential equilibrium and perfect equilibrium, since these play a key role in our new results. Halpern’s results involve the use of *nonstandard probability measures*, which take values in *non-Archimedean fields*. We briefly review these as well, and then state and prove the new characterizations of sequential equilibrium, quasi-perfect equilibrium, and perfect equilibrium in terms of common knowledge of rationality. For our results, we need to consider two types of rationality: *local rationality*, which considers only whether each player’s action is a best response at each information set (with everything else fixed), and *rationality*, which considers whether his whole strategy from that point on is a best response. This distinction seems critical when comparing perfect and quasi-perfect equilibrium (as already noted by van Damme [1984]); interestingly, it is not critical when it comes to sequential equilibrium. We compare our results to those of Asheim and Perea [2005], who provide a characterization of sequential equilibrium and quasi-perfect equilibrium for 2-player games in terms of common knowledge of rationality similar in spirit to ours. We conclude in Section 4 with a discussion of the use of common knowledge of rationality in characterizing solution concepts.

2 A review of earlier results

To explain our results, we briefly review the earlier results on characterizing solution concepts in strategic-form games terms of common knowledge (see [Dekel and Siniscalchi 2015] for a more comprehensive survey). We assume that the reader is familiar with standard solution concepts such as Nash equilibrium, correlated equilibrium, and rationalizability; see [Osborne and Rubinstein 1994] for a discussion. Let $\Gamma = (N, \mathcal{S}, (u_i)_{i \in N})$ be a finite strategic-form game, where $N = \{1, \dots, n\}$ is the set of players, $\mathcal{S} = \times_{i \in N} \mathcal{S}_i$ is a finite set of strategy profiles, and $u_i : \mathcal{S} \rightarrow \mathbb{R}$ is player i ’s utility function. For ease of exposition, we assume that $\mathcal{S}_i \cap \mathcal{S}_j = \emptyset$ for $i \neq j$.

Let a *model* of Γ be a tuple $M = (\Omega, \mathbf{s}, (\text{Pr}_i)_{i \in N})$, where Ω is a set of states of Γ , \mathbf{s} associates with each state $\omega \in \Omega$ a pure strategy profile $\mathbf{s}(\omega) \in \mathcal{S}$, and Pr_i is a

probability distribution on Ω , describing i 's initial beliefs.¹ Let $\mathbf{s}_i(\omega)$ denote player i 's strategy in the profile $\mathbf{s}(\omega)$, and let $\mathbf{s}_{-i}(\omega)$ denote the strategy profile consisting of the strategies of all players other than i .

For $S \in \mathcal{S}_i$, let $[S] = \{\omega \in \Omega : \mathbf{s}_i(\omega) = S\}$ be the set of states at which player i chooses strategy S . Similarly, let $[\vec{S}_{-i}] = \{\omega \in \Omega : \mathbf{s}_{-i}(\omega) = \vec{S}_{-i}\}$ and $[\vec{S}] = \{\omega \in \Omega : \mathbf{s}(\omega) = \vec{S}\}$. For simplicity, we assume that $[\vec{S}]$ is measurable for all strategy profiles \vec{S} , and that $\Pr_i([S_i]) > 0$ for all strategies $S_i \in \mathcal{S}_i$ and all players $i \in N$.

As usual, we say that a player is *rational* at state ω (in a model M of Γ) if his strategy at ω is a best response in Γ given his beliefs at ω . We view \Pr_i as i 's prior belief, intuitively, before i has been assigned or has chosen a strategy. We assume that i knows his strategy at ω , and that this is all that i learns in going from his prior knowledge to his knowledge at ω , so his beliefs at ω are the result of conditioning \Pr_i on $[\mathbf{s}_i(\omega)]$.² Given our assumption that $\Pr_i([\mathbf{s}_i(\omega)]) > 0$, the conditional probability $\Pr_i \mid [\mathbf{s}_i(\omega)]$ is well defined.

Note that we can view \Pr_i as inducing a probability \Pr_i^S on strategy profiles $\vec{S} \in \mathcal{S}$ by simply taking $\Pr_i^S(\vec{S}) = \Pr_i([\vec{S}])$; we similarly define $\Pr_i^S(S_i) = \Pr_i([S_i])$ and $\Pr_i^S(\vec{S}_{-i}) = \Pr_i([\vec{S}_{-i}])$. Let $\Pr_{i,\omega}^S = \Pr_i^S \mid \mathbf{s}_i(\omega)$. Intuitively, at state ω , player i knows his strategy $\mathbf{s}_i(\omega)$, so his distribution $\Pr_{i,\omega}^S$ on strategies at ω is the result of conditioning his prior distribution on strategies \Pr_i^S on this information.

Formally, i is *rational at ω* if, for all strategies $S \in \mathcal{S}_i$, we have that

$$\sum_{\vec{S}'_{-i} \in \mathcal{S}_{-i}} \Pr_{i,\omega}^S(\vec{S}'_{-i}) u_i(\mathbf{s}_i(\omega), \vec{S}'_{-i}) \geq \sum_{S'_{-i} \in \mathcal{S}_{-i}} \Pr_{i,\omega}^S(S'_{-i}) u_i(S, \vec{S}'_{-i}).$$

We say that player i is *rational in model M* if i is rational at every state ω in M . Finally, we say that *rationality is common knowledge in M* if all players are rational at every state of M . (Technically, our definition of rationality being common knowledge in M means that rationality is *universal* in M (i.e., true at *all* states in M), and thus, in particular, common knowledge at all states in M according to the standard definition of common knowledge at a state (cf., [Fagin et al. 1995]). While common knowledge of rationality at a state does not imply that rationality is universal in general, in the models that we focus on in this paper, the two notions coincide.)

With this background, we can state Aumann's [1987] characterization of Nash equilibrium. As usual, we can identify a mixed strategy profile $\vec{\sigma}$ in Γ with a distribution $\Pr_{\vec{\sigma}}$ on \mathcal{S} ; the distribution $\Pr_{\vec{\sigma}}$ can be viewed as a crossproduct $\times_{i \in N} \Pr_{\sigma_i}$ (where \Pr_{σ_i} is a distribution on \mathcal{S}_i).³ Let Σ_i denote the set of mixed strategies for player i .

¹For simplicity, we assume in this paper that Ω is finite, and all subsets of Ω are measurable.

²While this arguably is a reasonable assumption for strategic-form games, when we move to extensive-form games, agents will be able to learn more in the course of a game.

³We consistently use S , possibly subscripted, to denote a pure strategy, while σ , possibly subscripted or with a prime, denotes a mixed strategy.

Theorem 2.1 $\vec{\sigma}$ is a Nash equilibrium of Γ iff there exists a model $M = (\Omega, \mathbf{s}, (\Pr_i)_{i \in N})$ of Γ where rationality is common knowledge such that $\Pr_i = \Pr_j$ for all $i, j \in N$ and $\Pr_i^S = \Pr_{\vec{\sigma}}$ for all $i \in N$.

The fact that $\Pr_i = \Pr_j$ for all $i, j \in N$ means that there is a common prior. Because $\Pr_{\vec{\sigma}}$ has the form of a cross-product, the fact that $\Pr_i^S = \Pr_{\vec{\sigma}}$ means that i 's beliefs about other players' strategies is independent of the state; that is, $\Pr_i^S \mid \mathbf{s}_i(\omega)$ marginalized to \mathcal{S}_{-i} is independent of ω .⁴

Theorem 2.1 is actually a special case of Aumann's [1987] characterization of correlated equilibrium. Recall that we can think of a correlated equilibrium of Γ as a distribution η on \mathcal{S} . Intuitively, η is a correlated equilibrium if, when a mediator chooses a strategy profile \vec{S} according to η and tells each player i his component S_i of \vec{S} , then playing S_i is a best response for i . This intuition is formalized in Aumann's theorem:

Theorem 2.2 η is a correlated equilibrium of Γ iff there exists a model $M = (\Omega, \mathbf{s}, (\Pr_i)_{i \in N})$ of Γ where rationality is common knowledge such that $\Pr_i^S = \eta$ for all $i \in N$.

Theorems 2.1 and 2.2 show that the difference between correlated equilibrium and Nash equilibrium can be understood as saying that, with correlated equilibrium, the common prior does not have to be a cross-product, so that a player i 's beliefs may vary, for different choices of strategy. Of course, if the prior *is* a cross-product, then the correlated equilibrium is also a Nash equilibrium. With correlated equilibrium, as with Nash equilibrium, there is a common prior.

We complete the review of characterizations of solution concepts in strategic-form games in terms of common knowledge of rationality with the following characterization of correlated rationalizability (where a player can believe that other players' strategies are correlated), due to Brandenburger and Dekel [1987]:

Theorem 2.3 S_j is a (correlated) rationalizable strategy for player j in a game Γ iff there exists a model $M = (\Omega, \mathbf{s}, (\Pr_i)_{i \in N})$ of Γ where rationality is common knowledge and a state $\omega \in \Omega$ such that $\mathbf{s}_j(\omega) = S_j$.

Note that the characterization of rationalizability does not require the players to have a common prior.

⁴Aumann and Brandenburger [1995] show that common knowledge of rationality is not required for σ to be a Nash equilibrium. This is not a contradiction to Theorem 2.1, which simply says that σ is a Nash equilibrium iff there exists a model M describing the beliefs of the players where rationality is common knowledge. There may be other models where the players play σ and rationality is not common knowledge.

3 Characterizing sequential equilibrium and perfect equilibrium

Our goal is to characterize sequential equilibrium and perfect equilibrium in finite extensive-form games with perfect recall in terms of common knowledge of rationality. We assume that the reader is familiar with the standard definitions of extensive-form games of perfect (trembling hand) perfect equilibrium, quasi-perfect equilibrium, and sequential equilibrium. Our characterizations make essential use of non-epistemic characterizations of sequential and perfect equilibrium using nonstandard probability [Halpern 2009; Halpern 2013]. We briefly review these results here.

One of the issues that the definitions of sequential and perfect equilibrium need to deal with are probability zero events, specifically, those corresponding to information sets that are off the equilibrium path. Halpern [2009, 2013] presents a novel way to approach this issue in the context of games, by making use of nonstandard probability measures, which we now describe.

Non-Archimedean fields are fields that include the real numbers \mathbb{R} as a subfield, and also contain *infinitesimals*, which are positive numbers that are strictly smaller than any positive real number. The smallest such non-Archimedean field, commonly denoted $\mathbb{R}(\varepsilon)$, is the minimal field generated by adding to the reals a single infinitesimal, denoted by ε .⁵ $\mathbb{R}(\varepsilon)$ consists of all the rational expressions $f(\varepsilon)/g(\varepsilon)$, where $f(x)$ and $g(x)$ are polynomials with real coefficients and $g(0) \neq 0$. It is easy to see that this gives us a field that includes the reals and ε . We can place an order $<$ on the elements of $\mathbb{R}(\varepsilon)$ by taking $0 < \varepsilon < 1/r$ for all reals $r > 0$, and extending to all of $\mathbb{R}(\varepsilon)$ by assuming that standard properties of the reals (e.g., that $r^2 < r$ if $0 < r < 1$) continue to hold. Thus, $0 < \dots < \varepsilon^3 < \varepsilon^2 < \varepsilon$ holds, for all real numbers $r > 0$ we have that $1/\varepsilon > r$, and so on. (We can use formal division to identify $f(\varepsilon)/g(\varepsilon)$ with a power series of the form $a_0 + a_1\varepsilon + a_2\varepsilon^2 + \dots$; this suffices to guide how the order $<$ should be extended to quotients $f(\varepsilon)/g(\varepsilon)$.)

The field $\mathbb{R}(\varepsilon)$ does not suffice for our purposes. In this paper we will be interested in non-Archimedean fields \mathbb{R}^* that are *elementary extensions* of the standard reals. This means that \mathbb{R}^* is an ordered field that includes the real numbers, at least one infinitesimal ε , and is *elementarily equivalent to the field of real numbers*. The fact that \mathbb{R}^* and \mathbb{R} are elementarily equivalent means that every formula φ that can be expressed in first-order logic and uses the function symbols $+$ and \times (interpreted as addition and multiplication, respectively), and constant symbols \mathbf{r} standing for particular real numbers (the underlying language contains a constant symbol \mathbf{r} for each real number $r \in \mathbb{R}$) is true in F iff φ is true in \mathbb{R} . We call such a field a *normal non-Archimedean field*. Thus, for example, every odd-degree polynomial has a root in a normal non-Archimedean field \mathbb{R}^* since this fact is true in \mathbb{R} and can be expressed in first-order logic. Note that $\mathbb{R}(\varepsilon)$ is not a normal non-Archimedean field. For example, one property of the reals

⁵The construction of $\mathbb{R}(\varepsilon)$ apparently goes back to Robinson [1973].

expressible in first-order logic is that every positive number has a square root. However, ε does not have a square root in $\mathbb{R}(\varepsilon)$. For the results of this paper, we do not have to explicitly describe a normal non-Archimedean field; it suffices that one exists. The existence of normal non-Archimedean fields is well known, and follows from the fact that first-order logic is *compact*; see [Enderton 1972].⁶

Given a normal non-Archimedean field \mathbb{R}^* , we call the elements of \mathbb{R} the *standard* reals in \mathbb{R}^* , and those of $\mathbb{R}^* \setminus \mathbb{R}$ the *nonstandard* reals. A nonstandard real b is *finite* if $-r < b < r$ for some standard real $r > 0$. If $b \in \mathbb{R}^*$ is a finite nonstandard real, then $b = a + \varepsilon$, where a is the unique standard real number closest to b and ε is an infinitesimal. Formally, $a = \inf\{r \in \mathbb{R} : r > b\}$ and $\varepsilon = b - a$; it is easy to check that ε is indeed an infinitesimal. We call a the *standard part* of b , and denote it $st(b)$.

A nonstandard probability measure \Pr on Ω just assigns each event in Ω an element in $[0, 1]$ in some (fixed) non-Archimedean field \mathbb{R}^* . Note that $\Pr(\Omega) = 1$, just as with standard probability measures. We require \Pr to be finitely additive. Recall that, for the purposes of this paper, we restrict attention to finite state spaces Ω . This allows us to avoid having to define an analogue of countable additivity for nonstandard probability measures. Given a nonstandard probability measure ν , we can define the standard probability measure $st(\nu)$ by taking $st(\nu)(w) = st(\nu(w))$. Two possibly nonstandard distributions ν and ν' *differ infinitesimally* if $st(\nu) = st(\nu')$ (i.e., for all events E , the probabilities $\nu(E)$ and $\nu'(E)$ differ by at most an infinitesimal, so $st(\nu(E) - \nu'(E)) = 0$). If a nonstandard distribution assigns a positive (possibly infinitesimal) probability to every possible outcome in a game, then there is no technical problem in conditioning on such outcomes. Moreover, every standard probability measure differs infinitesimally from a nonstandard probability measure that assigns positive probabilities to all outcomes.

A *behavioral* strategy σ for player i in an extensive-form game associates with each information set I for player i a distribution $\sigma(I)$ over the actions that can be played at I . We allow $\sigma(I)$ to be a nonstandard probability distribution. We say that σ is *standard* if $\sigma(I)$ is standard for all information sets I for player i . Two behavioral strategy σ and σ' for player i *differ infinitesimally* if, for all information sets I for player i , the distributions $\sigma(I)$ and $\sigma'(I)$ differ infinitesimally. Two strategy profiles $\vec{\sigma}$ and $\vec{\sigma}'$ differ infinitesimally if σ_i and σ'_i differ infinitesimally for $i = 1, \dots, n$. We say that a behavioral strategy σ is *completely mixed* if it assigns positive (but possibly infinitesimal) probability to every action at every information set.

A behavioral strategy profile in an extensive-form game induces a probability on *terminal histories* of the game (i.e., histories that start at the root of the game tree and end at a leaf). Let Z_Γ be the set of terminal histories in a game Γ . (We omit explicit mention of the game Γ if it is clear from context or irrelevant.) Given a behavioral strategy profile $\vec{\sigma}$ for Γ , let $\Pr_{\vec{\sigma}}$ be the probability on terminal histories induced by $\vec{\sigma}$.

⁶There is a natural extension of $\mathbb{R}(\varepsilon)$ called $\mathbb{R}^*(\varepsilon)$ that is normal. As shown by [Halpern 2009; Halpern 2013], Theorems 3.3 and 3.5 could be strengthened to use $\mathbb{R}^*(\varepsilon)$ rather than an existentially quantified normal non-Archimedean field.

Thus, $\text{Pr}_{\vec{\sigma}}$ is a distribution on pure strategy profiles if $\vec{\sigma}$ is a mixed strategy profile, and a distribution on histories if $\vec{\sigma}$ is a behavioral strategy profile in an extensive-form game. We hope that the context will disambiguate the notation. Since we can identify a partial history with the terminal histories that extend it, $\text{Pr}_{\vec{\sigma}}(h)$ and $\text{Pr}_{\vec{\sigma}}(I)$ are well defined for a partial history h and an information set I . Recall that in an extensive-form game Γ , each player i 's utility function is defined on Z_Γ .

A *belief system* [Kreps and Wilson 1982] is a function μ that associates with each information set I a probability, denoted μ_I , on the histories in I . Given a behavioral strategy $\vec{\sigma}$ and a belief system μ in an extensive-form game Γ , let

$$\text{EU}_i((\vec{\sigma}, \mu) \mid I) = \sum_{h \in I} \sum_{z \in Z} \mu_I(h) \text{Pr}_{\vec{\sigma}}(z \mid h) u_i(z).$$

Thus, the expected utility of $(\vec{\sigma}, \mu)$ conditional on reaching I captures the expected payoff to player i if I is reached via the distribution $\vec{\sigma}$ and from that point on the game is played according to μ . Intuitively, this expected utility captures what i can expect to receive if i changes its strategy at information set I .

Finally, if $\vec{\sigma}$ is a completely-mixed behavioral strategy profile, let $\mu^{\vec{\sigma}}$ be the belief system determined by $\vec{\sigma}$ in the obvious way:

$$\mu_I^{\vec{\sigma}}(h) = \text{Pr}_{\vec{\sigma}}(h \mid I).$$

Definition 3.1 Fix a game Γ . Let I be an information set for player i , let $\vec{\sigma}'$ be a completely-mixed behavioral strategy profile, and let $\varepsilon \geq 0$. Then we say that σ_i is an ε -best response to $\vec{\sigma}'_{-i}$ for i conditional on having reached I using $\vec{\sigma}'$ if, for every strategy τ_i for player i , we have that

$$\text{EU}_i(((\sigma_i, \vec{\sigma}'_{-i}), \mu_I^{\vec{\sigma}'}) \mid I) \geq \text{EU}_i(((\tau_i, \vec{\sigma}'_{-i}), \mu_I^{\vec{\sigma}'}) \mid I) - \varepsilon. \quad (1)$$

The strategy σ_i is an ε -best response for i relative to $\vec{\sigma}'$ if σ_i is an ε -best response to $\vec{\sigma}'_{-i}$ for i conditional on having reached I using $\vec{\sigma}'$ for all information sets I for i . ■

Observe that in Equation 1 the probability of reaching I on both sides of the inequality depends only on $\vec{\sigma}'$ (via $\mu_I^{\vec{\sigma}'}$) and not on τ_i . Thus, τ_i only influences player i 's behavior after I has been reached.

Given an information set I for player i , let A_I be the set of actions available to i at histories in I .⁷ As usual, we take $\Delta(A_I)$ to be the set of probability measures on A_I . Note that if σ_i is a behavioral strategy for player i then, by definition, $\sigma_i(I) \in \Delta(A_I)$.

Definition 3.2 If $\varepsilon \geq 0$ and I is an information set for player i that is reached with positive probability by $\vec{\sigma}'$, then $a \in \Delta(A_I)$ is a *local ε -best response to $\vec{\sigma}'_{-i}$ for i conditional on having reached I using $\vec{\sigma}'$* if, for all $a' \in \Delta(A_I)$, we have that

$$\text{EU}_i(((\sigma'_i[I/a], \vec{\sigma}'_{-i}), \mu_I^{\vec{\sigma}'}) \mid I) \geq \text{EU}_i(((\sigma'_i[I/a'], \vec{\sigma}'_{-i}), \mu_I^{\vec{\sigma}'}) \mid I) - \varepsilon, \quad (2)$$

⁷As is standard, we assume that the same set of actions is available to i at all histories in I .

where $\sigma'_i[I/a']$ is the behavioral strategy that agrees with σ'_i except possibly at information set I , and $\sigma'_i[I/a'](I) = a'$. The strategy σ_i is a *local ε -best response for i relative to $\vec{\sigma}'$* if $\sigma_i(I)$ is a local ε -best response to $\vec{\sigma}'_{-i}$ for i conditional on having reached I using $\vec{\sigma}'$ for all information sets I for i . The strategy $\vec{\sigma}_i$ is a *(local) best response for i relative to $\vec{\sigma}'$* (resp., *(local) best response for i conditional on having reached I using $\vec{\sigma}'$*) if σ_i is a (local) 0-best response for i relative to $\vec{\sigma}'$ (resp., (local) 0-best response for i conditional on having reached I). ■

Thus, with local best responses, we consider the best *action* at an information set; with (non-local) best responses, we consider the best continuation strategy.

Halpern [2009, 2013] characterizes perfect equilibrium using non-Archimedean fields and local best responses as follows:

Theorem 3.3 *Let Γ be a finite extensive-form game with perfect recall. Then the (standard) behavioral strategy profile $\vec{\sigma} = (\sigma_1, \dots, \sigma_n)$ is a perfect equilibrium of Γ iff there exists a normal non-Archimedean field \mathbb{R}^* and a nonstandard completely-mixed behavioral strategy profile $\vec{\sigma}'$ with probabilities in \mathbb{R}^* that differs infinitesimally from $\vec{\sigma}$ such that, for each player $i = 1, \dots, n$ and each information set I of player i , $\sigma_i(I)$ is a local best response for i relative to $\vec{\sigma}'$.*

Roughly speaking, Theorem 3.3 shows that we can replace the sequence of strategies converging to $\vec{\sigma}$ considered in Selten's definition of perfect equilibrium by a single nonstandard completely-mixed strategy that is infinitesimally close to $\vec{\sigma}$. Considering a completely-mixed strategy guarantees that all information sets are reached with positive probability, and thus allows us to define best responses conditional on reaching an information set, for every information set.

We can obtain a characterization of quasi-perfect equilibrium by requiring that σ_i be a best response for i rather than a local best response [Halpern 2009; Halpern 2013].⁸ As we said earlier, the fact that the key difference between perfect equilibrium and quasi-perfect equilibrium is that local best responses were required for the former and best responses were required for the latter was already stressed by van Damme [1984] in his original definition of quasi-perfect equilibrium.

Theorem 3.4 [Halpern 2009; Halpern 2013] *Let Γ be a finite extensive-form game with perfect recall. Then the (standard) behavioral strategy profile $\vec{\sigma} = (\sigma_1, \dots, \sigma_n)$ is a quasi-perfect equilibrium of Γ iff there exists a normal non-Archimedean field \mathbb{R}^* and a nonstandard completely-mixed behavioral strategy profile $\vec{\sigma}'$ with probabilities in \mathbb{R}^* that differs infinitesimally from $\vec{\sigma}$ such that, for each player $i = 1, \dots, n$, the strategy σ_i is a best response for i relative to $\vec{\sigma}'$.*

⁸The characterization of perfect equilibrium given in [Halpern 2009] involved best responses. In [Halpern 2013], it was pointed out that this was incorrect; σ_i needed to be a *local* best response to get a characterization of perfect equilibrium, but taking it to be a best response gave a characterization of quasi-perfect equilibrium.

Finally, we can obtain a characterization of sequential equilibrium by requiring that σ_i be an ε -best response for i to $\vec{\sigma}'$ rather than a local best response as in Theorem 3.3, or a best response as in Theorem 3.4. It can be shown if ε is an infinitesimal, then there exists an infinitesimal ε' such that an ε -local best response relative to $\vec{\sigma}'$ is actually an ε' -best response (see Lemma 3.10), so, as we would expect, the requirement for sequential equilibrium is actually a weakening of the requirements for both perfect and quasi-perfect equilibrium.

Theorem 3.5 [Halpern 2009; Halpern 2013] *Let Γ be a finite extensive-form game with perfect recall. Then there exists a belief system μ such that the assessment $(\vec{\sigma}, \mu)$ is a sequential equilibrium of Γ iff there exist a normal non-Archimedean field \mathbb{R}^* , an infinitesimal $\varepsilon \in \mathbb{R}^*$, and a nonstandard completely-mixed behavioral strategy profile $\vec{\sigma}'$ with probabilities in \mathbb{R}^* that differs infinitesimally from $\vec{\sigma}$ such that σ_i is an ε -best response for i relative to $\vec{\sigma}'$, for each player $i = 1, \dots, n$.*

Our epistemic characterizations are based on Theorems 3.3, 3.4, and 3.5. Given a finite extensive-form game Γ , we take a model M of Γ to be a tuple $(\Omega, \mathbf{Z}, (\text{Pr}_i)_{i \in N})$ where, as before, Ω is a finite set of states and Pr_i is a (possibly nonstandard) probability distribution on Ω . Now \mathbf{Z} is a function that associates with each state $\omega \in \Omega$ a terminal history in Γ , denoted $\mathbf{Z}(\omega)$. The distribution Pr_i on states induces a distribution Pr_i^Z on terminal histories in the obvious way. A model $M = (\Omega, \mathbf{Z}, (\text{Pr}_i)_{i \in N})$ of the game Γ is *compatible with* a behavioral strategy profile $\vec{\sigma}$ if $\text{Pr}_1^Z = \dots = \text{Pr}_n^Z = \text{Pr}_{\vec{\sigma}}$.

We now define two notions of rationality, corresponding to the types of best response considered above: local best response and best response. To be consistent with the type of response considered, we call these *local rationality* and *rationality*. Both notions have been considered in the literature, although different terms have been used. Arieli and Aumann [2015] use the terms *action rationality* and *utility maximization* instead of “local rationality” and “rationality”.

Definition 3.6 Fix $\varepsilon > 0$ and a model M compatible with a completely-mixed strategy profile $\vec{\sigma}'$. Player i is *ε -locally rational at state ω* if, for each information set I for player i , if some history $h \in I$ is a prefix of $\mathbf{Z}(\omega)$, player i plays action a after h in $\mathbf{Z}(\omega)$, and $st(\sigma'_i(I)(a)) > 0$, then a is a local ε -best response to $\vec{\sigma}'_{-i}$ for i conditional on having reached I using $\vec{\sigma}'$. Player i is *locally rational at ω* if he is 0-locally rational at ω . Player i is *ε -rational at state ω* if, for each information set I for player i , if some history $h \in I$ is a prefix of $\mathbf{Z}(\omega)$, then $st(\sigma'_i)$ is an ε -best response to $\vec{\sigma}'_{-i}$ for i conditional on having reached I using $\vec{\sigma}'$. Player i is *rational at state ω* if he is 0-rational at ω . ■

Note that in the definition of local rationality at ω , we do not require that the action played by i at a prefix of $\mathbf{Z}(\omega)$ be a local best response if that action is played with only infinitesimal probability. Similarly, in the definition of rationality, we require $st(\sigma'_i)$ to be a best response, not σ'_i , since we are ultimately interested in $st(\sigma'_i)$. Also note that we

define rationality only in models that are compatible with a completely-mixed behavioral strategy profile. This ensures that the expected utility conditional on I is well defined for each information set I . We could, of course, try to define rationality more generally, but the extra work would not be relevant to the results of this paper.

We are now ready to formally capture perfect equilibrium in terms of common knowledge of rationality, using Theorem 3.3. Intuitively, the assumption that σ_i is a best response relative to the nonstandard $\vec{\sigma}$ is replaced by the assumption of common knowledge of rationality when players play $\vec{\sigma}'$.

Theorem 3.7 *Let Γ be a finite extensive-form game with perfect recall. Then $\vec{\sigma}$ is a perfect equilibrium of Γ iff there exist a normal non-Archimedean field \mathbb{R}^* , a nonstandard, completely-mixed strategy profile $\vec{\sigma}'$ that differs infinitesimally from $\vec{\sigma}$ with probabilities in \mathbb{R}^* , and a model $M = (\Omega, \mathbf{Z}, (\text{Pr}_i)_{i \in N})$ of Γ compatible with $\vec{\sigma}'$ where local rationality is common knowledge.*

Proof: Suppose that $\vec{\sigma}$ is a perfect equilibrium of Γ . Then, by Theorem 3.3, there exists a normal non-Archimedean field \mathbb{R}^* and a nonstandard completely-mixed strategy profile $\vec{\sigma}'$ with probabilities in \mathbb{R}^* that differs infinitesimally from $\vec{\sigma}$ such that, for each player i , the strategy σ_i is a local best response for i relative to $\vec{\sigma}'$. Let $M = (\Omega, \mathbf{Z}, (\text{Pr}_i)_{i \in N})$ be such that $\Omega = \{\omega_h : h \in Z_\Gamma\}$, $\mathbf{Z}(\omega_h) = h$, and $\text{Pr}_i(\omega_h) = \text{Pr}_{\vec{\sigma}'}(h)$, for $i = 1, \dots, n$. Clearly M is compatible with $\vec{\sigma}'$. We claim that it is common knowledge in M that all players are locally rational.

To see this, consider an arbitrary state $\omega_h \in \Omega$. Suppose that I is an information set for player i , $h' \in I$ is a prefix of h , the action played by i at h' in h is a , and $st(\sigma'_i(I)(a)) > 0$. Since σ_i is a local best response for i conditional on having reached I using $\vec{\sigma}'$, Equation (2) from Definition 3.2 (with $\varepsilon = 0$) implies that

$$\text{EU}_i((\sigma'_i[I/\sigma_i(I)], \vec{\sigma}'_{-i}), \mu_I^{\vec{\sigma}'}) \mid I \geq \text{EU}_i((\sigma_i[I/a'], \vec{\sigma}'_{-i}), \mu_I^{\vec{\sigma}'}) \mid I$$

for all $a' \in \Delta(A_I)$. It easily follows that

$$\text{EU}_i((\sigma'_i[I/a''], \vec{\sigma}'_{-i}), \mu_I^{\vec{\sigma}'}) \mid I \geq \text{EU}_i((\sigma_i[I/a'], \vec{\sigma}'_{-i}), \mu_I^{\vec{\sigma}'}) \mid I \quad (3)$$

for all actions $a' \in A_I$ and all actions a'' in the support of $\sigma_i(I)$. By assumption, σ'_i differs infinitesimally from σ_i . Hence, the fact that $st(\sigma'_i(I)(a)) > 0$ implies that $\sigma_i(I)(a) > 0$, so that the action a must be in the support of $\sigma_i(I)$. Therefore, (3) holds for $a' = a$, so i is rational at ω_h . We conclude that every player i is locally rational at all states $\omega \in \Omega$ and thus, by definition, it is common knowledge in M that the players are locally rational.

For the converse, fix $\vec{\sigma}$ and suppose that there exist \mathbb{R}^* , $\vec{\sigma}'$, and a model M as required by the theorem. For each information set I for player i , if $a \in A_I$ is in the support of $\sigma_i(I)$, then $st(\sigma'_i(I)(a)) > 0$. Since M is compatible with $\vec{\sigma}'$, there must exist some state ω in M with a prefix h of $\mathbf{Z}(\omega)$ in I such that i plays a after h in $\mathbf{Z}(\omega)$. Since i

is locally rational at ω , performing a must be a local best response for i conditional on having reached I using $\vec{\sigma}'$. Thus, $\sigma_i(I)$ must be a local best response for i conditional on having reached I using $\vec{\sigma}'$. Hence, by Theorem 3.3 we obtain that $\vec{\sigma}$ is a perfect equilibrium. ■

Perhaps not surprisingly, we obtain an analogue of Theorem 3.7 by replacing “local rationality” by “rationality”.

Theorem 3.8 *Let Γ be a finite extensive-form game with perfect recall. Then $\vec{\sigma}$ is a quasi-perfect equilibrium of Γ iff there exist a normal non-Archimedean field \mathbb{R}^* , a nonstandard, completely-mixed strategy profile $\vec{\sigma}'$ that differs infinitesimally from $\vec{\sigma}$ with probabilities in \mathbb{R}^* , and a model $M = (\Omega, \mathbf{Z}, (\Pr_i)_{i \in N})$ of Γ compatible with $\vec{\sigma}'$ where rationality is common knowledge.*

Proof: The proof is similar in spirit to that of Theorem 3.7, and simpler, so we leave details to the reader. ■

Interestingly, for sequential equilibrium, we can work with either ε -rationality or ε -local rationality.

Theorem 3.9 *Let Γ be a finite extensive-form game with perfect recall. The following are equivalent:*

- (a) *there exists a belief system μ such that the assessment $(\vec{\sigma}, \mu)$ is a sequential equilibrium of Γ ;*
- (b) *there exist a normal non-Archimedean field \mathbb{R}^* , a nonstandard, completely-mixed behavioral strategy profile $\vec{\sigma}'$ with probabilities in \mathbb{R}^* that differs infinitesimally from $\vec{\sigma}$, an infinitesimal $\varepsilon > 0$ in \mathbb{R}^* , and a model $M = (\Omega, \mathbf{Z}, (\Pr_i)_{i \in N})$ compatible with $\vec{\sigma}$ where ε -rationality is common knowledge;*
- (c) *there exist a normal non-Archimedean field \mathbb{R}^* , a nonstandard, completely-mixed strategy profile $\vec{\sigma}'$ with probabilities in \mathbb{R}^* that differs infinitesimally from $\vec{\sigma}$, an infinitesimal $\varepsilon > 0$ in \mathbb{R}^* , and a model $M = (\Omega, \mathbf{Z}, (\Pr_i)_{i \in N})$ compatible with $\vec{\sigma}$ where ε -local rationality is common knowledge.*

Proof: In light of Theorem 3.5, the equivalence of (a) and (b) is almost immediate. To see that (a) implies (c), suppose that $(\vec{\sigma}, \mu)$ is a sequential equilibrium. By Theorem 3.5, there exists a strategy profile $\vec{\sigma}'$ that differs infinitesimally from $\vec{\sigma}$ and an infinitesimal ε such that, for each player i , strategy σ_i is an ε -local best response relative to $\vec{\sigma}'_{-i}$. Construct M as in the proof of Theorem 3.7. Since σ_i differs infinitesimally from σ'_i , for each player i , there exists an infinitesimal ε'_i such that, for all information sets I for player i ,

$$\text{EU}_i(((\sigma_i, \vec{\sigma}'_{-i}), \mu_I^{\vec{\sigma}'})) \mid I) \geq \text{EU}_i(((\sigma_i, \vec{\sigma}'_{-i}), \mu_I^{\vec{\sigma}'})) \mid I) - \varepsilon'_i. \quad (4)$$

Let $\varepsilon' = \max_{i=1, \dots, n} \varepsilon'_i$ and let

$r = \min_{i=1, \dots, n} \{\sigma_i(I)(a) : I \text{ is an information set for } i \text{ and } a \text{ is in the support of } \sigma_i(I)\}$;

that is, r is the smallest positive probability assigned by a strategy σ_i , $i = 1, \dots, n$. Note that $\varepsilon' + \varepsilon + \varepsilon/r$ is an infinitesimal (since r is a standard rational). We claim that $(\varepsilon' + \varepsilon + \varepsilon/r)$ -local rationality is common knowledge in M .

To see this, fix a player i and a state ω in M , and let $h = \mathbf{Z}(\omega)$. Again, suppose that I is an information set for player i , $h' \in I$ is a prefix of h , the action played by i at h' in h is a , and $st(\sigma'_i(I))(a) > 0$. We want to show that

$$\text{EU}_i(((\sigma'_i[I/a], \vec{\sigma}'_{-i}), \mu_I^{\vec{\sigma}'}) \mid I) \geq \text{EU}_i(((\sigma'_i[I/a'], \vec{\sigma}'_{-i}), \mu_I^{\vec{\sigma}'}) \mid I) - (\varepsilon' + \varepsilon + \varepsilon/r) \quad (5)$$

for all actions $a' \in A_I$. First observe that, by choice of ε' , it easily follows from (4) that

$$\text{EU}_i(((\sigma'_i[I/a], \vec{\sigma}'_{-i}), \mu_I^{\vec{\sigma}'}) \mid I) \geq \text{EU}_i(((\sigma_i[I/a], \vec{\sigma}'_{-i}), \mu_I^{\vec{\sigma}'}) \mid I) - \varepsilon'. \quad (6)$$

Moreover, since σ_i is an ε -best response relative to $\vec{\sigma}'_{-i}$, for all actions $a' \in A_I$, we must have

$$\text{EU}_i(((\sigma_i, \vec{\sigma}'_{-i}), \mu_I^{\vec{\sigma}'}) \mid I) \geq \text{EU}_i(((\sigma'_i[I/a'], \vec{\sigma}'_{-i}), \mu_I^{\vec{\sigma}'}) \mid I) - \varepsilon. \quad (7)$$

Since (7) holds for each action a' in the support of $\sigma_i(I)$, we must have

$$\begin{aligned} & \text{EU}_i(((\sigma_i, \vec{\sigma}'_{-i}), \mu_I^{\vec{\sigma}'}) \mid I) \\ &= \sigma_i(I)(a) \text{EU}_i(((\sigma_i[I/a], \vec{\sigma}'_{-i}), \mu_I^{\vec{\sigma}'}) \mid I) + \sum_{\{a' : \sigma_i(I)(a') > 0, a' \neq a\}} \sigma_i(I)(a') \text{EU}_i(((\sigma_i[I/a'], \vec{\sigma}'_{-i}), \mu_I^{\vec{\sigma}'}) \mid I) \\ &\leq \sigma_i(I)(a) \text{EU}_i(((\sigma_i[I/a], \vec{\sigma}'_{-i}), \mu_I^{\vec{\sigma}'}) \mid I) + \sum_{\{a' : \sigma_i(I)(a') > 0, a' \neq a\}} \sigma_i(I)(a') (\text{EU}_i(((\sigma_i, \vec{\sigma}'_{-i}), \mu_I^{\vec{\sigma}'}) \mid I) + \varepsilon) \\ &= \sigma_i(I)(a) \text{EU}_i(((\sigma_i[I/a], \vec{\sigma}'_{-i}), \mu_I^{\vec{\sigma}'}) \mid I) + (1 - \sigma_i(I)(a)) (\text{EU}_i(((\sigma_i, \vec{\sigma}'_{-i}), \mu_I^{\vec{\sigma}'}) \mid I) + \varepsilon). \end{aligned}$$

A little algebraic manipulation now shows that

$$\begin{aligned} \text{EU}_i(((\sigma_i[I/a], \vec{\sigma}'_{-i}), \mu_I^{\vec{\sigma}'}) \mid I) &\geq \text{EU}_i(((\sigma_i, \vec{\sigma}'_{-i}), \mu_I^{\vec{\sigma}'}) \mid I) - \varepsilon(1 - \sigma_i(I)(a))/\sigma_i(I)(a) \\ &\geq \text{EU}_i(((\sigma_i, \vec{\sigma}'_{-i}), \mu_I^{\vec{\sigma}'}) \mid I) - \varepsilon/r. \end{aligned} \quad (8)$$

Equation (5) follows immediately from (6), (7), and (8). Thus, $(\varepsilon' + \varepsilon + \varepsilon/r)$ -local rationality is common knowledge in M . We have shown that (a) implies (c).

It remains to show that (c) implies (a). So suppose that there exists a field \mathbb{R}^* , a nonstandard strategy profile $\vec{\sigma}'$, an infinitesimal $\varepsilon > 0$ in \mathbb{R}^* , and a model M where ε -local rationality is common knowledge, as required for (c) to hold. It is almost immediate that σ_i is an ε -local best response for i relative to σ'_i . We want to show that there exists some infinitesimal ε'' , possibly different from ε , such that σ_i is an ε'' -best response for i relative to σ'_i , for each player i . The result then follows from Theorem 3.5.

To do this, we need some preliminary definitions. In a finite extensive-form game Γ with perfect recall, for each player i , we can define a partial order \succ_i on player i 's information sets such that $I \succ_i I'$ if, for every history $h \in I$, there is a prefix h' of h in I' .

Thus, $I \succ_i I'$ if I is below (i.e., appears later than) I' in the game tree. We define the *height* of an information set I for player i , denoted by $height(I)$, inductively as follows; $height(I) = 1$ if I is a maximal set for player i , that is, there is no information set I' such that $I' \succ_i I$. If I is not maximal, then $height(I) = \max\{height(\hat{I}) + 1 : \hat{I} \succ_i I\}$. Since Γ is a finite game, $height(I)$ is well defined. Indeed, the size of the game ensures that there is a finite bound d such that $height(I) \leq d$ for all information sets in the game. For ε' defined just before Equation (4), we can now prove the following result:

Lemma 3.10 σ_i is a $d(\varepsilon + \varepsilon')$ -best response for i relative to $\vec{\sigma}'$ in Γ .

Proof: For all information sets I of player i , we show by induction on $k = height(I)$ that σ_i is a $k(\varepsilon + \varepsilon')$ -best response to $\vec{\sigma}'_i$ conditional on having reached I using $\vec{\sigma}'$. So fix an arbitrary player i , and let I be an information set for i . If I is maximal, then $height(I) = 1$. By assumption, σ_i is a local ε -best response to $\vec{\sigma}'_{-i}$ conditional on having reached I using $\vec{\sigma}'$, so the base case of the induction holds. Now suppose that $height(I) = k > 1$ and that the claim holds for all I' such that $height(I') < k$.

By choice of ε' , we have by Equation (4) that

$$EU_i(((\sigma_i, \vec{\sigma}'_{-i}), \mu_I^{\vec{\sigma}'}) \mid I) \geq EU_i(((\sigma'_i[I/\sigma_i(I)], \vec{\sigma}'_{-i}), \mu_I^{\vec{\sigma}'}) \mid I) - \varepsilon'. \quad (9)$$

Let τ_i be an arbitrary strategy for player i . By assumption, σ_i is a local ε -best response relative to $\vec{\sigma}'_{-i}$, so

$$EU_i(((\sigma'_i[I/\sigma_i(I)], \vec{\sigma}'_{-i}), \mu_I^{\vec{\sigma}'}) \mid I) \geq EU_i(((\sigma'_i[I/\tau_i(I)], \vec{\sigma}'_{-i}), \mu_I^{\vec{\sigma}'}) \mid I) - \varepsilon. \quad (10)$$

Let $\mathcal{I} = \{I_1, \dots, I_m\}$ be the information sets for player i that immediately succeed I in Γ (i.e., for each $I_j \in \mathcal{I}$, $I_j \succeq I$ and there is no information set I' such that $I_j \succ_i I' \succ_i I$) and can be reached by starting at a history in I and playing $\tau(I)$. By the inductive hypothesis, σ_i is a $(k-1)(\varepsilon + \varepsilon')$ -best response to $\vec{\sigma}'_{-i}$ at each information set $I' \in \mathcal{I}$, so player i 's utility is at most $(k-1)(\varepsilon + \varepsilon')$ worse if he plays σ_i rather than τ_i at each $I' \in \mathcal{I}$. It easily follows that

$$EU_i(((\sigma_i[I/\tau_i(I)], \vec{\sigma}'_{-i}), \mu_I^{\vec{\sigma}'}) \mid I) \geq EU_i((\tau_i, \vec{\sigma}'_{-i}), \mu_I^{\vec{\sigma}'}) \mid I) - (k-1)(\varepsilon + \varepsilon'). \quad (11)$$

Putting together (9), (10), and (11), we obtain that

$$EU_i(((\sigma_i, \vec{\sigma}'_{-i}), \mu_I^{\vec{\sigma}'}) \mid I) \geq EU_i((\tau_i, \vec{\sigma}'_{-i}), \mu_I^{\vec{\sigma}'}) \mid I) - k(\varepsilon + \varepsilon').$$

Since $height(I) \leq d$ for each information set I in Γ , it follows that σ_i is an $d(\varepsilon + \varepsilon')$ -best response for i relative to $\vec{\sigma}'$, for each player $i = 1, \dots, n$. This completes the proof of the lemma. ■

Clearly $\varepsilon'' = d(\varepsilon + \varepsilon')$ is an infinitesimal, so by Theorem 3.5, it follows that there exists a belief system μ such that the assessment $(\vec{\sigma}, \mu)$ is a sequential equilibrium of Γ , as desired. ■

It follows from Theorem 3.9 that Theorem 3.5 can be generalized to use either ε -rationality or ε -local rationality. Each of the choices gives a characterization of sequential rationality.

It is interesting to compare our results to those of Asheim and Perea [2005]. As mentioned, they provide epistemic characterizations of sequential equilibrium and quasi-perfect equilibrium in 2-player games in terms of rationality. Their notion of rationality is essentially equivalent to ours; since they do not use local rationality, it is perhaps not surprising that they do not deal with perfect equilibrium, which seems to require it.

To obtain their results, Asheim and Perea represent uncertainty using a generalization of LPSs (lexicographic probability sequences) [Blume, Brandenburger, and Dekel 1991a; Blume, Brandenburger, and Dekel 1991b] that they call *systems of conditional lexicographic probabilities* (SCLPs). An LPS is a sequence (Pr_0, \dots, Pr_k) of probability measures on a measure space (S, \mathcal{F}) . Roughly speaking, we can identify such a sequence with the nonstandard probability measure $(1 - \epsilon - \dots - \epsilon^k) Pr_0 + \epsilon Pr_1 + \dots + \epsilon^k Pr_k$ on (S, \mathcal{F}) . Indeed, it has been shown that LPSs and nonstandard probability spaces (NPSs) are essentially equivalent in finite spaces [Blume, Brandenburger, and Dekel 1991a; Halpern 2010]. However, it is not hard to show that SCLPs can capture some situations that cannot be captured by NPSs. Roughly speaking, this is because SCLPs do not necessarily satisfy an analogue of the chain rule of probability ($\Pr(A | B) \times \Pr(B | C) = \Pr(A | C)$ if $A \subseteq B \subseteq C$), which does hold for NPSs.⁹ (Of course, we might view such situations as unreasonable.) It would be interesting to investigate whether our results could be obtained with some variant of LPSs or CPSs (conditional probability spaces).

Another relatively minor difference between our result and that of Asheim and Perea is that they work with what they call *common certain belief* rather than with common knowledge, where certain belief of E is defined relative to a model characterized by an LPS (μ_1, \dots, μ_k) if $\mu_j(E) = 1$ for $j = 1, \dots, k$. Although Asheim and Perea's theorems are stated in terms of *mutual* certain belief of rationality rather than common certain belief, where mutual certain belief holds if both of the players have certain belief of rationality, they also require mutual certain belief of each player's type; in their setting, this implies common certain belief of rationality.

Finally, in their characterization of quasi-perfect equilibrium, Asheim and Perea also require common certain belief of *caution*, which, roughly speaking, in our language, says that players should prefer a strategy that is a best response to one that is an ε -best response, even for an infinitesimal ε . Dropping caution when moving from quasi-perfect equilibrium to sequential equilibrium in Asheim and Perea's framework corresponds to moving from rationality to ε -rationality in our framework.

⁹There is no notion of multiplication in SCLPs, so this statement is not quite accurate. Nevertheless, consequences of the chain rule, such as that $\mu(A | B) = \mu(A' | B)$ implies $\mu(A | C) = \mu(A' | C)$ do not hold for SCLPs.

4 Discussion

Theorems 3.7, 3.8, and 3.9 illustrate the role that common knowledge of rationality plays in perfect equilibrium, quasi-perfect equilibrium, and sequential equilibrium. Comparing Theorem 2.1 to Theorem 3.7, note that for $\vec{\sigma}$ to be a perfect equilibrium, Theorem 3.7 requires players to *always* be rational; that is, for every information set I that a player i can reach in the game, i must be rational conditional on reaching I . Since Theorem 2.1 considers only normal-form games, the requirement that players always be rational has no bite. But we could prove an analogue of Theorem 2.1 for Nash equilibrium in extensive-form games, and again it would suffice to have rationality *ex ante*, rather than conditional on reaching each information set. The other key difference between Theorems 2.1 and 3.7 is that in Theorem 3.7, rather than taking the probability on histories in M to be determined by $\vec{\sigma}$, it is determined by $\vec{\sigma}'$, a completely-mixed nonstandard strategy that differs infinitesimally from $\vec{\sigma}$. Note that there are many strategies that differ infinitesimally from $\vec{\sigma}$. The exact choice of $\vec{\sigma}'$ has only an infinitesimal impact on i 's beliefs at information sets I that are on the equilibrium path; but for information sets I off the equilibrium path, the choice of $\vec{\sigma}'$ completely determines i 's beliefs; different choices can result in quite different beliefs.

The distinction between Theorems 3.7 and 3.9 highlights one way of thinking about the difference between perfect equilibrium and sequential equilibrium. For perfect equilibrium, it has to be common knowledge that players are always rational; for sequential equilibrium, it suffices to have common knowledge that players are always ε -rational for an infinitesimal $\varepsilon > 0$. The distinction between Theorems 3.7 and 3.8 brings out the point that van Damme already stressed in the definition of quasi-perfect equilibrium: the difference between local best responses and best responses. We find it of interest that this distinction does not play a role in sequential equilibrium.

Our results complement Aumann's earlier epistemic characterizations of Nash and of correlated equilibria. The general picture obtained is that all of these solution concepts can be characterized in terms of common knowledge of rationality; the differences between the characterizations depend on what we assume about the prior probability, whether rationality holds at all information sets or just at the beginning, and whether we consider rationality or ε -rationality. As we show in related work [Halpern and Moses 2007], as a consequence of this observation, it follows that all these solution concepts can be embodied in terms of a single *knowledge-based program* [Fagin et al. 1995; Fagin et al. 1997], which essentially says that player i should perform action a if she believes both that she plans to perform a and that playing a is optimal for her in the sense of being a best response. This is, arguably, the essence of rationality. In the case of each of the equilibrium notions that we have discussed, for the corresponding notions of rationality and best response, if it is common knowledge that everyone is following this knowledge-based program, then rationality is common knowledge.

Can other standard solution concepts be characterized this way? It is straightforward to state and prove an analogue of Theorem 2.1 for Bayes-Nash equilibrium. Now the

state space in the model would include each player’s type. If we define rationality and best responses in terms of minimax regret, rather than in terms of maximizing expected utility, Boutilier and Hyafil [2004] define a notion of *minimax-regret* equilibrium that can be captured in terms of common knowledge of rationality. Similarly, Aghassi and Bertsimas [2006] define rationality in terms of maximin (i.e., maximizing the worst-case utility) and use that to define what they call *maximin equilibria*. Again, we can prove an analogue of Theorem 2.1 for this solution concept.

Perhaps more interesting is the solution concept of *iterated admissibility*, also known as *iterated deletion of weakly dominated strategies*. Brandenburger, Friedenberg, and Keisler [2008] provide an epistemic characterization of iterated admissibility (i.e., iterated deletion of weakly dominated strategies) where uncertainty is represented using LPSs (lexicographic probability sequences). They define a notion of belief (which they call *assumption*) appropriate for their setting, and show that strategies that survive k rounds of iterated deletion are ones that are played in states where there is k th-order mutual belief in rationality; that is, everyone assumes that everyone assumes \dots ($k - 1$ times) that everyone is rational. However, they prove only that their characterization of iterated admissibility holds in particularly rich structures called *complete* structures, where all types are possible. However, more recently, Halpern and Pass [2009] provide a characterization that is closer to the spirit of Theorem 2.1. The key new feature is that instead of just requiring that everyone is rational, and that everyone knows that everyone is rational, and that everyone knows that everyone knows \dots , they require that *all* everyone knows is that everyone is rational, and that all everyone knows is that all everyone knows is that everyone is rational, and so on. In this claim, the statement that *all agent i knows is φ* is true at a state ω if, not only is it the case that φ is true at all states that i considers possible at ω (which is what is required for i to know φ at ω), but it is also the case that i assigns ψ positive probability for each formula ψ consistent with φ . Thus, we capture “all i knows is φ ” by requiring that i considers any situation compatible with φ possible. In the specific case of iterated admissibility, this means that i considers possible (i.e., assigns positive probability to) all strategies compatible with rationality. As shown by Halpern and Pass [2009], a strategy survives k rounds of iterated deletion iff it is played at a state in a structure where all everyone knows is that all everyone knows \dots (k times) that everyone is rational. This result does not require the restriction to complete structures.

Now consider *extensive-form rationalizability* (EFR) [Pearce 1984], an extension of rationalizability that seems appropriate for extensive-form games [Halpern and Pass 2009]. Battigalli and Siniscalchi [2002] provide an epistemic characterization of EFR using a notion of *strong belief*; these are beliefs that are maintained unless evidence shows that the beliefs are inconsistent. For example, if player 1 has a strong belief of player 2’s rationality, then whatever moves player 2 makes, player 1 will revise her beliefs and, in particular, her beliefs about player 2’s beliefs, in such a way that she continues to believe that player 2 is rational (so that she believes that player 2 is making a best response to his beliefs), unless it is inconsistent for her to believe that player 2 is rational. Battigalli

and Siniscalchi characterize EFR in terms of common strong belief of rationality. Specifically, they show that a strategy satisfies EFR iff it is played in a complete structure. Again, using “all i knows” would allow us to give an epistemic characterization of EFR in the spirit of the theorems in this paper without the restriction to complete structures [Halpern and Pass 2009].¹⁰

To summarize, the notion of common knowledge of rationality seems deeply embedded in many game-theoretic solution concepts. While not all solution concepts can be given epistemic characterizations in terms of some variant of common knowledge of rationality (one counterexample is the notion of *iterated regret minimization* [Halpern and Pass 2012]), the results of this paper and of others mentioned in the previous discussion show that many of the most popular solution concepts do admit such a characterization.

Acknowledgements

We would like to thank the anonymous reviewers for their detailed reading of the paper and useful comments that helped improve the paper. In particular, we thank an anonymous reviewer for encouraging us to compare our results to those of Asheim and Perea [2005].

References

- Aghassi, M. and D. Bertsimas (2006). Robust game theory. *Mathematical Programming, Series B* 107(1–2), 231–273.
- Arieli, I. and R. J. Aumann (2015). The logic of backward induction. *Journal of Economic Theory* 159, 443–464.
- Asheim, G. B. and A. Perea (2005). Sequential and quasi-perfect rationalizability in extensive games. *Games and Economic Behavior* 53, 15–42.
- Aumann, R. J. (1987). Correlated equilibrium as an expression of Bayesian rationality. *Econometrica* 55, 1–18.
- Aumann, R. J. and A. Brandenburger (1995). Epistemic conditions for Nash equilibrium. *Econometrica* 63(5), 1161–1180.
- Battigalli, P. and M. Siniscalchi (2002). Strong beliefs and forward-induction reasoning. *Journal of Economic Theory* 106, 356–391.
- Blume, L., A. Brandenburger, and E. Dekel (1991a). Lexicographic probabilities and choice under uncertainty. *Econometrica* 59(1), 61–79.
- Blume, L., A. Brandenburger, and E. Dekel (1991b). Lexicographic probabilities and equilibrium refinements. *Econometrica* 59(1), 81–98.

¹⁰Perea [2012] also provides epistemic characterizations of iterated admissibility and EFR that do not require complete type structures.

- Brandenburger, A. and E. Dekel (1987). Rationalizability and correlated equilibria. *Econometrica* 55, 1391–1402.
- Brandenburger, A., A. Friedenberg, and J. Keisler (2008). Admissibility in games. *Econometrica* 76(2), 307–352.
- Dekel, E. and M. Siniscalchi (2015). Epistemic game theory. In H. P. Young and S. Zamir (Eds.), *Handbook of Game Theory with Economic Applications, Volume 4*, pp. 619–702. North-Holland.
- Enderton, H. B. (1972). *A Mathematical Introduction to Logic*. New York: Academic Press.
- Fagin, R., J. Y. Halpern, Y. Moses, and M. Y. Vardi (1995). *Reasoning About Knowledge*. Cambridge, Mass.: MIT Press. A slightly revised paperback version was published in 2003.
- Fagin, R., J. Y. Halpern, Y. Moses, and M. Y. Vardi (1997). Knowledge-based programs. *Distributed Computing* 10(4), 199–225.
- Halpern, J. Y. (2009). A nonstandard characterization of sequential equilibrium, perfect equilibrium, and proper equilibrium. *International Journal of Game Theory* 38(1), 37–50.
- Halpern, J. Y. (2010). Lexicographic probability, conditional probability, and nonstandard probability. *Games and Economic Behavior* 68(1), 155–179.
- Halpern, J. Y. (2013). A nonstandard characterization of sequential equilibrium, perfect equilibrium, and proper equilibrium: Erratum. Unpublished manuscript.
- Halpern, J. Y. and Y. Moses (2007). Characterizing solution concepts in games using knowledge-based programs. In *Proc. Twentieth International Joint Conference on Artificial Intelligence (IJCAI '07)*, pp. 1300–1307.
- Halpern, J. Y. and R. Pass (2009). A logical characterization of iterated admissibility and extensive-form rationalizability. A preliminary version, with the title "A logical characterization of iterated admissibility", appears in *Proc. Twelfth Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*, 2009, pp. 146–155.
- Halpern, J. Y. and R. Pass (2012). Iterated regret minimization: a new solution concept. *Games and Economic Behavior* 74(1), 194–207.
- Hyafil, N. and C. Boutilier (2004). Regret minimizing equilibria and mechanisms for games with strict type uncertainty. In *Proc. Twentieth Conference on Uncertainty in Artificial Intelligence (UAI 2004)*, pp. 268–277.
- Kreps, D. M. and R. B. Wilson (1982). Sequential equilibria. *Econometrica* 50, 863–894.
- Osborne, M. J. and A. Rubinstein (1994). *A Course in Game Theory*. Cambridge, Mass.: MIT Press.

- Pearce, D. G. (1984). Rationalizable strategic behavior and the problem of perfection. *Econometrica* 52(4), 1029–1050.
- Perea, A. (2012). *Epistemic Game Theory*. Cambridge, U.K.: Cambridge University Press.
- 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.
- van Damme, E. (1984). A relationship between perfect equilibria in extensive form games and proper equilibria in normal form games. *International Journal of Game Theory* 13, 1–13.