Journal of Mathematical Economics 72 (2017) 82–87

Contents lists available at ScienceDirect

Journal of Mathematical Economics

journal homepage: www.elsevier.com/locate/jmateco

Interim correlated rationalizability in infinite games ☆

Jonathan Weinstein [a,*], Muhamet Yildiz [b]

[a] Washington University in St. Louis, United States
[b] MIT, United States

ARTICLE INFO

Article history:
Received 27 May 2016
Received in revised form 3 July 2017
Accepted 5 July 2017
Available online 20 July 2017

Keywords:
Rationalizability
Incomplete information

ABSTRACT

In a Bayesian game, assume that the type space is a complete, separable metric space, the action space is a compact metric space, and the payoff functions are continuous. We show that the iterative and fixed-point definitions of interim correlated rationalizability (ICR) coincide, and ICR is non-empty-valued and upper hemicontinuous. This extends the finite-game results of Dekel et al. (2007), who introduced ICR. Our result applies, for instance, to discounted infinite-horizon dynamic games.

1. Introduction

Interim correlated rationalizability (henceforth ICR) has emerged as the main notion of rationalizability for Bayesian games. Among other reasons, it has the following two desirable properties: Firstly, it is upper hemicontinuous in types, i.e., one cannot obtain a substantially new ICR solution by perturbing a type. Secondly, two distinct definitions of ICR coincide: The fixed-point definition states that ICR is the weakest solution concept with a specific best-response property, and this is the definition that gives ICR its epistemological meaning, as a characterization of actions that are consistent with common knowledge of rationality. Under an alternate definition, ICR is computed by iteratively eliminating the actions that are never a best response (type by type), and this iterative definition is often more amenable for analysis.

The above properties were originally proven for games with finite action spaces and finite sets of payoff parameters (by Dekel et al., 2006, 2007, who introduced ICR). However, in many important games these sets are quite large. For example, in the infinitely repeated prisoners' dilemma game, the set of outcomes is uncountable. Hence, the set of all possible payoff functions (with payoffs in [0, 1]) is the unit cube with uncountably many dimensions. Therefore, to analyze the infinitely repeated prisoners' dilemma game under payoff uncertainty without restricting the payoffs, we would need to consider a space of underlying payoff parameters larger than the continuum. The action space is also uncountably infinite, because the action space consists of all mappings that specify whether one cooperates or defects at each of infinitely many histories.

In this note, we establish the aforementioned two properties of ICR in greater generality so that it can be applied to commonly-used games with large action and type spaces. Specifically, we establish that ICR is upper hemicontinuous and its iterative and fixed-point definitions coincide in games where (i) the spaces of payoff parameters and types are complete, separable metric spaces, (ii) the space of actions is a compact metric space, and (iii) the belief and payoff functions are continuous.

An immediate consequence of our result is that ICR is upper hemicontinuous when types are endowed with the product topology of belief hierarchies. That is, one cannot obtain a substantially distinct ICR solution for a type by perturbing the type's higher-order beliefs (possibly by considering another type space). In that sense, ICR predictions are robust to higher-order uncertainty. Interestingly, ICR is unique in that regard: any strict refinement of ICR, such as Bayesian Nash equilibrium, fails to be upper hemicontinuous as a correspondence from belief hierarchies (Weinstein and Yildiz, 2007).

When the action space is a compact metric space, the set of all continuous payoff functions is a complete separable metric space under the uniform metric. Then, the universal type space is also complete and separable metric space. Therefore, we obtain an upper-hemicontinuity result for ICR on universal type space under the same conditions for individual best response: the action space is a complete separable metric space and it is common knowledge that the utility functions are continuous.

In Section 6, we apply our results to infinite-horizon games with finitely many moves at every stage, where the relevant notion of continuity is continuity at infinity. As examples, we show that ICR is upper hemicontinuous in discounted repeated games with

☆ In earlier versions of the paper, the results were confined to compactly metrizable type spaces. We thank an anonymous referee for detailed comments and for providing the main arguments that extended our proofs to the type spaces that are completely metrizable and separable.

* Corresponding author.
E-mail address: j.weinstein@wustl.edu (J. Weinstein).

http://dx.doi.org/10.1016/j.jmateco.2017.07.002
0304-4068/

uniformly bounded stage-game payoffs, even if discount factors are unknown and stage payoffs are unknown, history-dependent or even stochastic. We applied our results here to general infinite-horizon games, including repeated games, in Weinstein and Yildiz (2013).

Our note relates to the existing literature as follows. The usual solution concepts, including Nash equilibrium and rationalizability under complete information (Bernheim, 1984; Pearce, 1984), are upper hemicontinuous with respect to the parameters of the game. However, as we mentioned above, most incomplete-information versions of these concepts are not upper hemicontinuous with respect to higher-order beliefs, and we are not aware of any positive result on this matter other than the result of Dekel, Fudenberg and Morris that we extend here. On the other hand, there is a sizeable literature on the equivalence of the iterative and fixed-point definitions of rationalizability under complete information. The equivalence holds when the action spaces are compact metric spaces and the payoff functions are continuous (Tan and Werlang, 1988), and we extend this result to Bayesian games (using ICR as the notion of rationalizability).[1] Under complete information, Lipman (1994) shows that the equivalence may fail in games with countably infinite, discrete action spaces, because reaching the fixed point may require iterations through a transfinite ordinal; Arieli (2010) discusses the same issue in continuous games, in which the payoff functions are continuous and the strategy spaces are complete separable metric spaces, finding that the iterations may need to reach the first uncountable ordinal (but not further). These results show that our continuity and compactness assumptions are not superfluous for the equivalence result.[2]

2. Basic definitions

Given any topological space X, we will write $\Delta(X)$ for the space of Borel probability measures on X and endow it with the weak topology. We will also endow all product spaces with their product topology. We write $f : X \rightrightarrows Y$ to mean that f is a correspondence from X to Y, i.e. an arbitrary subset of $X \times Y$ with the convention that $f(x)$ denotes $\{y : (x, y) \in f\}$. We consider the defining subset of $X \times Y$ to be the *graph* of the correspondence and denote this $G(f)$ for clarity.

The correspondence f is said to have the *closed-graph property* if $G(f)$ is closed. In general, the closed-graph property is weaker than upperhemicontinuity, but the two concepts coincide when Y is compact. When Y is compact, a correspondence $f : X \rightrightarrows Y$ is said to be *upper hemicontinuous* (UHC) if f has the closed-graph property.

Consider a Bayesian game $(N, A, u, \Theta, T, \kappa)$ in normal form where

- N is a set of players,
- $A = \times_{i\in N} A_i$ is a compact metric space of action profiles,[3]
- Θ is a complete, separable metric space of payoff parameters,
- $u_i : \Theta \times A \to [0, 1]$ is a continuous payoff function for each player $i \in N$,
- $T = \times_{i\in N} T_i$ is a complete, separable metric space of type profiles, and
- $\kappa_{t_i} \in \Delta(\Theta \times T_{-i})$ is the *interim* belief of type t_i on the payoff functions and the other players' types. The map $t_i \mapsto \kappa_{t_i}$ is assumed to be continuous.

In the definition of a Bayesian game above, we made some topological assumptions that we will maintain throughout. These assumptions are satisfied in many applications, including discounted dynamic games, as discussed in Section 6.

For each player i, define the best-response correspondence $BR_i : \Delta(\Theta \times A_{-i}) \rightrightarrows A_i$ by setting

$$BR_i(\mu) = \arg\max_{a_i \in A_i} \int u_i(\theta, a_i, a_{-i})\, d\mu(\theta, a_{-i})$$

for each $\mu \in \Delta(\Theta \times A_{-i})$, where arg max is the set of maximizers. The best-response correspondence is extended to $\Delta(\Theta \times T_{-i} \times A_{-i})$ by $BR_i(\mu) = BR_i\left(\mathrm{marg}_{\Theta\times A_{-i}}\mu\right)$, where marg_X takes the marginal with respect to X. Since A_i is compact and u_i is continuous and bounded, BR_i is always non-empty. Moreover,

Fact 1. *BR_i is UHC.*

This fact is a version of Berge's Maximum Theorem suitable for our purposes; see the Appendix for a straightforward proof.

We next present the two definitions of interim correlated rationalizability (ICR). The definitions may differ in general infinite games, but will coincide under the present conditions. The first definition is given by iterated elimination of strictly dominated actions that are never a weak best response, as follows:

Define the family of correspondences $S_i^m : T_i \rightrightarrows A_i, i \in N, m \in \mathbb{N}$, iteratively, by setting $S_i^0(t_i) = A_i$ and for each $m > 0$ and $t_i \in T_i$, let $a_i \in S_i^m(t_i)$ if and only if there exists $\mu \in \Delta(\Theta \times T_{-i} \times A_{-i})$ such that

$$a_i \in BR_i(\mu), \tag{2.1}$$

$$\kappa_{t_i} = \mathrm{marg}_{\Theta\times T_{-i}}\mu, \tag{2.2}$$

$$\mu\left(G\left(S_{-i}^{m-1}\right)\right) = 1. \tag{2.3}$$

Here, the first condition requires that a_i is a best response to belief μ; the second condition requires that μ is a belief of type t_i (coherence), and the last condition requires that the other players play according to S_{-i}^{m-1} under μ. The set $G\left(S_{-i}^{m-1}\right)$ is closed and therefore measurable, by Proposition 1. Here, the domain of S_{-i}^{m-1} is taken to be $\Theta \times T_{-i}$, where $\theta \in \Theta$ does not affect $S_{-i}^{m-1}(\theta, t_{-i})$. We consider such larger domains whenever it is convenient. The limiting correspondence $S^\infty : T \rightrightarrows A_i$ is defined by

$$S^\infty(t) = \bigcap_{m\geq 0} S^m(t). \tag{2.4}$$

The correspondence S^∞ is our first definition of ICR.

The second definition is given as a fixed-point property. A solution concept $f : T \rightrightarrows A$ is said to have the *best-response property* if for every $t_i \in T_i$ and $a_i \in f(t_i)$, there exists $\mu \in \Delta(\Theta \times T_{-i} \times A_{-i})$ such that

$$a_i \in BR_i(\mu), \tag{2.5}$$

$$\kappa_{t_i} = \mathrm{marg}_{\Theta\times T_{-i}}\mu, \tag{2.6}$$

$$\mu(G(f_{-i})) = 1. \tag{2.7}$$

By convention, we take (2.7) to mean that μ assigns probability 1 to a measurable subset of $G(f_{-i})$. Interim correlated rationalizability is defined as the largest correspondence $R : T \rightrightarrows A$ with the best-response property. Note that, since the best-response property is

[1] The equivalence was also discussed, for a more narrow class of games, in Bernheim (1984), and in Dufwenberg and Stegeman (2002). This extension to Bayesian games was also covered by Ely and Peski (2006), using interim independent rationalizability as the solution concept.

[2] Also related is Chen et al. (2007) who form an iterative definition which is independent of the order of elimination, by allowing dominance by previously eliminated strategies; standard concepts may be order-dependent without the presence of regularity assumptions, as noted by Dufwenberg and Stegeman (2002).

[3] For any family $(X_i)_{i\in N}$ of sets and any family of functions $f_i : X_i \to Y_i, i \in N$, we write $X = \times_{i\in N} X_i$ and $X_{-i} = \times_{j\in N\setminus\{i\}} X_j$ and define functions $f : X \to Y$ and $f_{-i} : X_{-i} \to Y_{-i}$ by setting $f_{-i}(x_{-i}) = \left(f_j(x_j)\right)_{j\in N\setminus\{i\}}$.

closed under coordinate-wise union, R is well-defined, i.e.,

$$R_i(t_i) = \bigcup_{f:T\rightrightarrows A \text{ with best-response property}} f_i(t_i) \quad (t_i \in T_i)$$

has the best-response property. Note also that the only difference from the iterative definition of S^∞ is that the definition of R requires that the other players play according to R as well, while S^m requires only that they play according to S^{m-1}. Consequently, R is a stronger solution concept.

Fact 2. *For any t_i and $m \le \infty$, $R_i(t_i) \subseteq S_i^m(t_i)$.*

Proof. It suffices to show this for any finite m. The statement is true for $m = 0$, by definition. Towards an induction, assume that it is true for $m - 1$. For any $a_i \in R_i(t_i)$, there exists μ with (2.5)–(2.7). But since $R_{-i} \subseteq S_{-i}^{m-1}$ by the inductive hypothesis, μ also satisfies (2.1)–(2.3), showing that $a_i \in S_i^m(t_i)$. □

When the two definitions differ, the fixed-point definition should be taken as the definition of rationalizability because it characterizes the strategic implications of common knowledge of rationality. However, under the present assumptions, the definitions are equivalent, as we show in the next section.

3. Upperhemicontinuity

In this section, we will show that S^∞ is UHC and non-empty, and further that it coincides with the fixed-point definition.

Proposition 1. *For every $m \le \infty$, S^m is UHC and non-empty.*

Proof. For each finite $m \in \mathbb{N}$, we will show that S^m has the closed-graph property, i.e., $G(S^m)$ is closed. This further implies that

$$G(S^\infty) = \bigcap_{m\ge 0} G(S^m)$$

is also closed. Since A is compact, this is indeed the desired result: S^m is UHC for each $m \le \infty$.

Clearly, $G(S^0) = T \times A$ is closed. Towards an induction, assume that $G(S_{-i}^{m-1})$ is closed for some $i \in N$ and $m \in \mathbb{N}$. Take a sequence $(t_{i,k}, a_{i,k}) \in G(S_i^m)$ with limit (t_i, a_i). For each k, since $(t_{i,k}, a_{i,k}) \in G(S_i^m)$, there exists $\mu_k \in \Delta(\Theta \times T_{-i} \times A_{-i})$ such that

$$a_{i,k} \in BR_i(\mu_k), \tag{3.1}$$

$$\kappa_{t_{i,k}} = \text{marg}_{\Theta\times T_{-i}}\mu_k, \tag{3.2}$$

$$\mu_k\left(G\left(S_{-i}^{m-1}\right)\right) = 1. \tag{3.3}$$

Now, because κ is continuous, $\kappa_{t_{i,k}}$ is convergent, hence (as a set) it is relatively compact, and Prohorov's theorem tells us that it is tight. That is, for each $\varepsilon > 0$ there is a set $K_\varepsilon \subset \Theta \times T_{-i}$ such that $\kappa_{t_{i,k}}(K_\varepsilon) > 1 - \varepsilon$ for all k. Because A_{-i} is compact, each $K_\varepsilon \times A_{-i}$ is also compact, and by (3.2), $\mu_k(K_\varepsilon \times A_{-i}) = \kappa_{t_{i,k}}(K_\varepsilon) > 1 - \varepsilon$. Thus μ_k is tight. By Prohorov's Theorem, μ_k is relatively compact, and hence has a convergent subsequence. Focus now on this subsequence and call the limit μ. We will show that μ satisfies the conditions (2.1)–(2.3), showing that $a_i \in S_i^m(t_i)$, as desired.

Firstly, since $\Theta \times T_{-i} \times A_{-i}$ is endowed with product topology, the projection mapping is continuous and hence

$$\kappa_{t_{i,k}} = \text{marg}_{\Theta\times T_{-i}}\mu_k \to \text{marg}_{\Theta\times T_{-i}}\mu,$$

where the equality is by (3.2). Since $\kappa_{t_{i,k}} \to \kappa_{t_i}$ by continuity of κ, this further implies that

$$\text{marg}_{\Theta\times T_{-i}}\mu = \kappa_{t_i},$$

proving (2.2). Secondly, since BR_i is UHC by Fact 1, (3.1) implies that $a_i \in BR_i(\mu)$, proving (2.1). Finally, since $G(S_{-i}^{m-1})$ is closed (by the inductive hypothesis) and $\mu_k \to \mu$, by Portmanteau Theorem,

$$\mu\left(G\left(S_{-i}^{m-1}\right)\right) \ge \limsup \mu_k\left(G\left(S_{-i}^{m-1}\right)\right) = 1,$$

where the last equality is by (3.3). This proves (2.3) and completes the proof.

Non-emptiness holds for each finite m because compactness implies that best-reply sets are non-empty, and then follows for S^∞ by the finite intersection principle. □

By Fact 2, R is always contained in S^∞. The next result establishes the reverse inclusion.

Proposition 2. $S^\infty = R$.

Proof. By Fact 2, it suffices to show that $S^\infty \subseteq R$. We will prove this by showing that S^∞ has the best-response property. To this end, take any $t_i \in T_i$ and $a_i \in S_i^\infty(t_i)$. Now, since $a_i \in S_i^m(t_i)$ for each m, there exists a sequence $\mu_m \in \Delta(\Theta \times T_{-i} \times A_{-i})$ such that

$$a_i \in BR_i(\mu_m),$$

$$\kappa_{t_i} = \text{marg}_{\Theta\times T_{-i}}\mu_m,$$

$$\mu_m\left(G\left(S_{-i}^{m-1}\right)\right) = 1.$$

As in the proof of Proposition 1, μ_m has a limit μ with $a_i \in BR_i(\mu)$ and $\kappa_{t_i} = \text{marg}_{\Theta\times T_{-i}}\mu$. We will also show that $\mu\left(G\left(S_{-i}^\infty\right)\right) = 1$, showing that S^∞ has the best-response property. Now, for any $m \in \mathbb{N}$, since $\mu_k\left(G\left(S_{-i}^m\right)\right) = 1$ for every $k > m$ and $G\left(S_{-i}^m\right)$ is closed, by the Portmanteau Theorem, $\mu\left(G\left(S_{-i}^m\right)\right) = 1$. Since $G\left(S_{-i}^m\right) \downarrow G\left(S_{-i}^\infty\right)$, this implies that

$$\mu\left(G\left(S_{-i}^\infty\right)\right) = \lim_m \mu\left(G\left(S_{-i}^m\right)\right) = 1,$$

completing the proof. □

Combining the two propositions above, we come to the main conclusion in this section:

Proposition 3. *R is UHC and non-empty.*

We must emphasize that the above results extend the results of Dekel et al. (2006, 2007), dropping their finiteness assumptions on Θ and A. In particular, Proposition 1 extends Theorem 2 of Dekel et al. (2006), and Proposition 2 extends Proposition 4 of Dekel et al. (2007). Our proofs are also similar to theirs.

4. Upperhemicontinuity in belief hierarchies

In Bayesian games, a type is meant to represent a belief hierarchy in the interim stage. A central issue in game theory is whether the solution is robust to small specification errors in modeling the belief hierarchies using types. Such robustness is formalized by upperhemicontinuity in belief hierarchies. In this section, we show that ICR is UHC in belief hierarchies, by embedding the type spaces in the universal type space and applying Proposition 3 to the universal type space.

Fix Θ, A, and utility functions $u_i : \Theta \times A \to [0, 1]$, and vary the type spaces (T, κ). We maintain the assumptions on Θ, A, u_i, and κ, but we allow T to have any topology. A solution concept Σ is defined as a mapping that yields a correspondence $\Sigma(\cdot|T, \kappa) : T \rightrightarrows A$ on each type space (T, κ). When we refer to ICR as a solution concept, we use the obvious notation $R(\cdot|T, \kappa)$ and $S^\infty(\cdot|T, \kappa)$, with the same meanings as in Section 2.

The type spaces are continuously embedded in the *universal type space* (T^*, κ^*) using the following belief-hierarchy mapping (see Mertens and Zamir, 1985; Brandenburger and Dekel, 1993). For any type space (T, κ) and type $t_i \in T_i$, belief hierarchy

$$h_i(t_i|T,\kappa) = \left(h_i^1(t_i|T,\kappa), h_i^2(t_i|T,\kappa), h_i^3(t_i|T,\kappa), \ldots\right)$$

is computed inductively by setting

$$h_i^1(t_i|T,\kappa) = \kappa_{t_i} \circ \rho_\Theta^{-1}$$

and

$$h_i^k(t_i|T,\kappa) = \kappa_{t_i} \circ \left(\rho_\Theta, h_{-i}^1(\cdot|T,\kappa), \ldots, h_{-i}^{k-1}(\cdot|T,\kappa)\right)^{-1}$$

where $\rho_\Theta : \Theta \times T_{-i} \to \Theta$ is the projection mapping. Here, $h_i^1(t_i|T,\kappa)$ is the first-order belief of type t_i—about the payoff parameter θ, and $h_i^k(t_i|T,\kappa)$ is the kth-order belief of type t_i—about the payoff parameter θ and first $k-1$ orders of beliefs. Each level of beliefs is given the weak-star topology. The universal type space T^* consists of all belief hierarchies obtained as above and is endowed with product topology. Since Θ is a complete and separable metric space, the universal type space T^* is also a complete and separable metric space, and κ^* is a continuous function of belief hierarchies. Note also that h is continuous.[4] Hence, all of our assumptions hold for the Bayesian game $(N, A, u, \Theta, T^*, \kappa^*)$, to which we apply our previous results. Our first result establishes that ICR only depends on the belief hierarchies and its iterative and fixed-point definitions coincide (although T need not be compact).

Proposition 4. *For any type space (T, κ), we have $R(\cdot|T,\kappa) = S^\infty(\cdot|T,\kappa)$, and*

$$S_i^\infty(t_i|T,\kappa) = S_i^\infty\left(h_i(t_i|T,\kappa)|T^*,\kappa^*\right) \qquad (\forall i \in N, t_i \in T_i). \quad (4.1)$$

Proof. Fix any (T, κ). It is straightforward to obtain (4.1) from the definitions; see Proposition 1 of Dekel et al. (2007) for a derivation under finiteness. Since T^* is a complete and separable metric space, by Proposition 2, we also have

$$R\left(\cdot|T^*,\kappa^*\right) = S^\infty\left(\cdot|T^*,\kappa^*\right). \quad (4.2)$$

This further implies that $R(\cdot|T,\kappa) = S^\infty(\cdot|T,\kappa)$. To see this, take any $t_i \in T_i$ and any $a_i \in S_i^\infty(t_i|T,\kappa)$. By (4.1), $a_i \in S_i^\infty(h_i(t_i|T,\kappa)|T^*,\kappa^*)$. However, by (4.2), $S^\infty(\cdot|T^*,\kappa^*)$ has the best-response property. Hence, a_i is a best response to a belief $\mu^* \in \Delta\left(\Theta \times T_{-i}^* \times A_{-i}\right)$ with (i) $\text{marg}_{\Theta \times T_{-i}^*}\mu^* = \kappa^*_{h_i(t_i|T,\kappa)}$ and (ii) $\mu^*\left(a_{-i} \in S_{-i}^\infty\left(t_{-i}^*|T^*,\kappa^*\right)\right) = 1$. Now we will define a belief $\mu \in \Delta(\Theta \times T_{-i} \times A_{-i})$ with similar properties to μ^*. Let $M^*(\cdot|\cdot)$ be the conditional probability system induced by μ^*, where $M^*(\cdot|\theta, t_{-i}^*) \in \Delta(A_{-i})$ for each $\left(\theta, t_{-i}^*\right) \in \Theta \times T_{-i}^*$, and let

$$\mu(E) = \int M^*(E_{(\theta,t_{-i})}|\theta, h_{-i}(t_{-i}))d\kappa_{t_i}(\theta, t_{-i})$$

where $E_{(\theta,t_{-i})} = \{a_{-i} : (\theta, t_{-i}, a_{-i}) \in E\}$ is a cross-section of E. Since $\text{marg}_{\Theta\times A_{-i}}\mu = \text{marg}_{\Theta\times A_{-i}}\mu^*$, a_i is also a best response to belief μ. Moreover, by (i), $\text{marg}_{\Theta\times T_{-i}}\mu = \kappa_{t_i}$ (by definition of κ^*), and by (ii) and (4.1), $\mu\left(a_{-i} \in S_{-i}^\infty(t_{-i}|T,\kappa)\right) = 1$. □

Proposition 4 establishes that ICR solution set does not depend on the way one models the belief hierarchies. In contrast, many solution concepts, such as Bayesian Nash equilibrium and "interim independent rationalizability", fail (4.1), and the solution set depends on the way one models the belief hierarchies. Indeed, if the original type space has redundant types, some of the equilibria may disappear when the type space is embedded in the universal type space; see for example Ely and Peski (2006). There then exists some type t_i in some type space (T, κ) such that $\Sigma(t_i|T,\kappa) \neq \Sigma(h_i(t_i|T,\kappa)|T^*,\kappa^*)$.[5]

We now turn to the main concept of this section: upperhemicontinuity in belief hierarchies. We use the following notion of convergence among types.

Definition 1. A sequence of types t_i^m from type spaces (T^m, κ^m) is said to *converge in belief hierarchies* to a type t_i from a type space (T, κ), denoted by $t_i^m \to t_i$, if

$$h_i^k\left(t_i^m|T^m,\kappa^m\right) \to h_i^k(t_i|T,\kappa) \quad \forall k. \quad (4.3)$$

Here, we use the product topology on belief hierarchies, which reflects the point of view of an observer who can only access to finite orders of beliefs (see Weinstein and Yildiz, 2007 for a detailed discussion). Our notion of convergence among types reflects the same view.

Definition 2. A solution concept Σ is said to be *upper hemicontinuous in belief hierarchies* if

$$\left[a_i^m \in \Sigma\left(t_i^m|T^m,\kappa^m\right)\forall m\right] \Longrightarrow a_i \in \Sigma(t_i|T,\kappa)$$

for every sequence of actions a_i^m with limit a_i and for every sequence of types t_i^m from type spaces (T^m, κ^m) that converges in belief hierarchies to a type t_i from a type space (T, κ).

Observe that when (4.1) fails, the solution concept cannot be UHC in belief hierarchies, no matter what topology one uses on the belief hierarchies. More broadly, if $h_i(t_i|T,\kappa) = h_i\left(t_i'|T',\kappa'\right)$ while $a_i \notin \Sigma(t_i|T,\kappa)$ for some $a_i \in \Sigma\left(t_i'|T',\kappa'\right)$, then the condition for UHC in belief hierarchies fails for the constant sequence $\left(a_i^m, t_i^m\right) = \left(a_i, t_i'\right)$. In particular, Bayesian Nash equilibrium and interim independent rationalizability are not UHC in belief hierarchies. Indeed, as we mentioned in the introduction, no strict refinement of ICR can be UHC in belief hierarchies. ICR turns out to be UHC in belief hierarchies, as Dekel, Fudenberg, and Morris show for finite A and Θ. Applying Proposition 3 to the universal type space (T^*, κ^*), the next result shows that ICR remains UHC in belief hierarchies in our more general setup.

Proposition 5. *R is upper hemicontinuous in belief hierarchies.*

Proof. Consider a sequence of actions a_i^m with limit a_i and a sequence of types $t_i^m \in T_i^m$ from type spaces (T^m, κ^m) with limit $t_i \in T_i$ from a type space (T, κ) in the sense of (4.3). Assume that for each m, $a_i^m \in R_i\left(t_i^m|T^m,\kappa^m\right)$. Then, for each m, we have

$$\begin{aligned} a_i^m \in R_i\left(t_i^m|T^m,\kappa^m\right) &= S_i^\infty\left(t_i^m|T^m,\kappa^m\right) \\ &= S_i^\infty\left(h_i\left(t_i^m|T^m,\kappa^m\right)|T^*,\kappa^*\right), \end{aligned}$$

where the first equality is by Proposition 2 and the second equality is by (4.1). Since $S^\infty(\cdot|T^*,\kappa^*)$ is UHC (by Proposition 1) and $h_i\left(t_i^m|T^m,\kappa^m\right) \to h_i(t_i|T,\kappa)$, we have

$$a_i \in S_i^\infty\left(h_i(t_i|T,\kappa)|T^*,\kappa^*\right) = S_i^\infty(t_i|T,\kappa) = R_i(t_i|T,\kappa),$$

where the equalities are by Proposition 4. □

5. Sufficient conditions

In this section we will present simple sufficient conditions under which our results immediately apply. We will assume throughout that the set A of action profiles is a compact metric space and the utility functions are continuous in a for each fixed θ. Then, we

4 See Mertens and Zamir (1985) for details. They assume that the type spaces are compact, but compactness is not needed for continuity of h.

5 These solution concepts may still have selections that are invariant to the way one models the hierarchies (Yildiz, 2015).

will show that all of our results apply, provided that we measure the distance between payoff functions θ according to the uniform metric (5.1). Of course, the uniform topology is generally finer than the product topology, so that the upperhemicontinuity result does not apply to some sequences of types that would converge if we used the product topology on payoff functions. However, if we restrict θ to a class of equicontinuous payoff functions, the uniform and product topologies coincide, so that the upperhemicontinuity result applies under the product topology as well.

Consider an n-player game in which A is a compact metric space, and let Θ^* be the set of all continuous payoff functions $\theta : A \to [0, 1]^n$. Endow Θ^* with the uniform metric d_u, defined by

$$d_u\left(\theta, \theta'\right) = \sup_{a \in A} \left\| \theta(a) - \theta'(a) \right\|. \tag{5.1}$$

It is well-known that (Θ^*, d_u) is a complete separable metric space. Let (T^*, κ^*) be the (Θ^*, d_u)-based universal type space. By Brandenburger and Dekel (1993), T^* is a complete and separable metric space under the product topology on belief hierarchies (as in previous sections). Finally, note that, under this formulation, the utility functions $u_i : \Theta \times A \to [0, 1]$ are defined by $u_i(\theta, a) = \theta_i(a)$, and u_i is a continuous function under the above metric. Therefore, all of our assumptions are satisfied, and ICR is UHC on $(N, A, u, \Theta^*, T^*, \kappa^*)$. In particular, R is upper hemicontinuous in belief hierarchies. Assuming the uniform metric on payoff functions, this obtains upperhemicontinuity of ICR under the same conditions needed for the best response correspondence: the action space is a compact metric space and the payoffs are continuous.

One may need to use coarser topologies on the payoff functions $\theta : A \to [0, 1]^n$, such as the topology of pointwise convergence. Under such topologies, upperhemicontinuity is a more stringent condition as it applies to larger sets of convergent sequences. It is worth mentioning that this distinction vanishes under certain restrictions on the payoff functions. In particular, suppose Θ is an *equicontinuous* set of payoff functions $\theta : A \to [0, 1]^n$.[6] Then, the Arzela–Ascoli theorem tells us that pointwise convergence implies uniform convergence. That is, the uniform and product topology coincide, so that the sup norm actually metrizes the product topology. Thus, by the previous paragraph, our results hold for the Θ-based universal type space under the product topology on Θ.[7] In particular, R is upper hemicontinuous in belief hierarchies under pointwise convergence, as long as it is common knowledge that the payoff functions are in an equicontinuous family Θ.

6. Examples

Now, we will exhibit standard classes of dynamic games in which the assumptions are satisfied. We will focus on infinite-horizon dynamic games, with discrete time, in which the set of possible moves is finite at any given history. In these games, the histories are of the form $h = (b_1, b_2, \ldots)$.[8] We measure the distance between any two histories h and h' by

$$d(h, h') = 2^{-L(h,h')},$$

where $L\left(h, h'\right)$ is the date of earliest discrepancy between h and h'. This endows the set of histories with the product topology. A plan of action a_i for a player i is a mapping that maps each history at which player i is to move to an available move at that history. Once again, we measure the distance between any two action plans a_i and a_i' by $d\left(a_i, a_i'\right) = 2^{-l}$ where l is the length of the shortest $h = (b_1, \ldots, b_l)$ with $a_i(h) \neq a_i\left(h'\right)$. Since the players can move only at histories of finite length, the set $A = A_1 \times \cdots \times A_n$ is a compact metric space. Let Z be the set of terminal histories. Each action profile $a \in A$ leads to an outcome $z(a) \in Z$. The outcome function $a \mapsto z(a)$ is continuous.

As in the previous section we consider payoff functions $\theta : Z \to [0, 1]^n$ and define the utility functions $u_i : \Theta \times A \to [0, 1]$ by $u_i(\theta, a) = \theta_i(z(a))$. A function $\theta : Z \to [0, 1]^n$ is said to be *continuous at infinity* if for every $\varepsilon > 0$ and $z \in Z$, there exists $\bar{L}$ such that $\left|\theta_i(z) - \theta_i\left(z'\right)\right| < \varepsilon$ whenever $L\left(z, z'\right) > \bar{L}$, i.e., whenever z and z' agree in the first $\bar{L}$ entries (Fudenberg and Levine, 1983). This is equivalent to continuity with respect to the metric on Z defined above. Let Θ^* be the set of all payoff functions $\theta : Z \to [0, 1]^n$ that are continuous at infinity. Apply the uniform metric d_u to Θ^*. Then, as in the previous section, all of our assumptions are satisfied by the (Θ^*, d_u)-based universal type space (T^*, κ^*). Therefore, ICR is UHC on $(N, A, u, \Theta^*, T^*, \kappa^*)$.

Once again, under an equicontinuity assumption, this leads to upperhemicontinuity under pointwise-convergence. A set Θ of functions $\theta : Z \to [0, 1]^n$ is said to be *equicontinuous at infinity* if for every $\varepsilon > 0$ and $z \in Z$, there exists $\bar{L}$ such that for every $\theta \in \Theta$, $\left|\theta_i(z) - \theta_i\left(z'\right)\right| < \varepsilon$ whenever $L\left(z, z'\right) > \bar{L}$. This strengthens the usual notion of continuity at infinity by picking $\bar{L}$ simultaneously for all $\theta \in \Theta$. As before, if it is common knowledge that the set of payoff functions is restricted to a family Θ which is equicontinuous at infinity, ICR is upper hemicontinuous in belief hierarchies which inherit the pointwise notion of convergence on payoff functions.

The next section shows how our results apply to the specific class of discounted repeated games.

6.1. Repeated games

Consider an infinitely repeated game with perfect-monitoring and with finite set $S = S_1 \times \cdots \times S_n$ of strategy profiles in the stage game. We first focus on the standard case of time-separable and stationary payoffs. Under this standard formulation, the payoff of a player i at any given history $h = \left(s^0, s^1, \ldots\right)$ is

$$U(\delta_i, g_i, h) = (1 - \delta_i) \sum_{l=0}^{\infty} \delta_i^l g_i(s^l)$$

where δ_i is the discount factor of player i and $g_i \in [0, 1]^S$ is his stage game payoff function. The function U is known, but δ_i and g_i may or may not be known. We write $\Theta^* = (0, 1)^n \times [0, 1]^{N \times S}$ for these possibly unknown parameters $\theta = (\delta_1, \ldots, \delta_n, g_1, \ldots, g_n)$ and endow it with the usual Euclidean topology. With appropriately chosen metric on $(0, 1)$ and standard metric on $[0, 1]$, Θ is separable and complete.[9] Let (T^*, κ^*) be the Θ-based universal type space. By Brandenburger and Dekel (1993), T^* is also a complete separable metric space. The payoff function of a player i is given by

$$u_i(\theta, a) = U(\delta_i, g_i, z(a)).$$

Since U and $a \mapsto z(a)$ are continuous, u_i is also continuous. Moreover, belief function κ^* is continuous with respect to the types in the universal type space T^*. Therefore, the game $(N, A, u, \Theta, T^*, \kappa^*)$ satisfies all of our assumptions, and ICR is UHC and can be computed by iterated strict dominance on this game. In other words, our results here show that, in a repeated game, ICR is UHC and can be computed by iterated strict dominance whenever the repeated game structure is common knowledge.

---

6 A family $F \subseteq Y^X$ of functions $f : X \to Y$ is said to be *equicontinuous* if for every x and every ϵ, there exists a δ such that for all x' with $d(x', x) < \delta$, and all $f \in F$, $d(f(x), f(x')) < \epsilon$.

7 Note that we do not need to worry that Θ is complete, i.e. a closed subset of Θ^*, because we can apply our result to the Θ^*-based universal type space, and continuity is inherited on the subset with common belief in the restriction to Θ.

8 See Osborne and Rubinstein (1994) and Weinstein and Yildiz (2013) for more details on the framework.

9 The metric on $(0, 1)$ is given by $d\left(x, x'\right) = \left|f(x) - f\left(x'\right)\right|$ where $f(x) = \log(x/(1 - x))$.

In some applications, players' payoffs are not time-separable or stationary. Consider a repeated game as in the previous paragraph, but allow non-separable and non-stationary payoffs. In particular, let the payoff of a player i at any given history $h = \left(s^0, s^1, \ldots\right)$ be

$$\theta_i\left(h\right) = U\left(\delta_i, g_i, h\right) \equiv \left(1-\delta_i\right) \sum_{l=0}^{\infty} \delta_i^l g_i^l \left(s^1, \ldots, s^l\right) \tag{6.1}$$

where δ_i is the discount factor of player i and $g_i = \left(g_i^1, g_i^2 \ldots\right)$ is a sequence of stage game payoff functions $g_i^l : S^l \to [0, 1]$, where the stage game payoff at any period l can depend on any action that has been taken up to period l. Separable (but non-stationary) payoffs are included in this class as the special case in which each g_i^l depends only in s^l. Write

$$\Theta^* = (0, 1)^n \times \prod_{l=1}^{\infty} [0, 1]^{N \times S^l} \tag{6.2}$$

for the set of parameters $\theta = (\delta_1, \ldots, \delta_n, g_1, \ldots, g_n)$. Under the product topology, Θ^* is completely metrizable and separable. Therefore, our general results apply: ICR is UHC and can be computed by iterated strict dominance.

6.2. Stochastic games

In stochastic games, the stage-game payoff functions are allowed to depend on a stochastic state. If the state and the realized stage-game payoffs are not observable, then these games can be encoded as a special case of the general repeated games described in Eqs. (6.1) and (6.2). Indeed, θ describes payoffs on all possible paths of stage-game payoffs, allowing the stochastic process and induced payoffs to be encoded as a belief about θ in the universal type space. This trick will not encompass stochastic games with any monitoring of the stochastic state, because in our setup of Section 6.1 the players only receive information about actions taken, not about θ. We now show how to model stochastic games with public monitoring of the state. The class of games modeled here is neither a subset nor a superset of that in the previous section.[10]

For any finite set Y of states, consider a (controlled) stochastic process $y_l|_{l=0}^{\infty}$ where the distribution of $y_l \in Y$ is a function of past states and the moves in the previous period. That is, the process is specified by providing, for each $l \geq 0$, a conditional distribution $\rho_l\left(y_l | y_0, \ldots, y_{l-1}, s^{l-1}\right)$. The stage-game payoff of player i at l is a function g_i^l of $\left(s_i^l, y_{l+1}\right)$, i.e. only directly of his action and the next state, though y_{l+1} may depend in turn on all players' moves including at time l. Players can only observe current and past states $y_0, \ldots, y_l$ and recall their previous moves, i.e., a player i moves at histories of the form $\left(y_0, s_i^0, \ldots, y_{l-1}, s_i^{l-1}, y_l\right)$. The set of parameters, which encode both payoffs and the initial and conditional distributions of y, can be written as

$$\Theta^* = \prod_{i \in N} \left((0, 1) \times \prod_{l=1}^{\infty} [0, 1]^{S_i \times Y} \right) \times \Delta(Y) \times \prod_{l=1}^{\infty} \Delta(Y)^{Y^l \times S}.$$

Once again, under the product topology, Θ^* is completely metrizable and separable. Therefore, our general results apply: ICR is UHC – with respect to hierarchies of beliefs over both payoffs and the laws governing y – and it can be computed by iterated strict dominance.

[10] It does not encompass the arbitrary non-stationary games of the previous section because we only allow finitely many states, so that arbitrary dependence on arbitrarily long histories would not be possible.

Appendix. Proof of Fact 1

It follows from the following.

Lemma 1. *Let (S, d_S) and (X, d_X) be metric spaces with S compact and X complete and separable, and let $u : S \times X \to \mathbb{R}$ be continuous. Then,*

$$B_i\left(\mu\right) = \arg\max_{s \in S} \int u\left(s, x\right) d\mu\left(x\right)$$

is upper hemicontinuous in μ.

Proof. By the Maximum Theorem, it suffices to show that

$$U\left(s, \mu\right) = \int u\left(s, x\right) d\mu\left(x\right)$$

is continuous. Indeed, it is the composition of the following maps, which are each continuous by standard results:

$$\left(s, \mu\right) \mapsto \left(\delta_s, \mu\right)$$

where δ_s is the unit mass at s;

$$\left(\delta_s, \mu\right) \mapsto \delta_s \times \mu$$

the product measure; and

$$\delta_s \times \mu \mapsto \int_{S \times X} u(s', x) d(\delta_s \times \mu)(s', x)$$

which is continuous by definition of the weak topology. □

References

Arieli, I., 2010. Rationalizability in continuous games. J. Math. Econom. 46, 912–924.
Bernheim, D., 1984. Rationalizable strategic behavior. Econometrica 52, 1007–1028.
Brandenburger, A., Dekel, E., 1993. Hierarchies of beliefs and common knowledge. J. Econom. Theory 59, 189–198.
Chen, Y., Long, N., Luo, X., 2007. Iterated strict dominance in general games. Games Econom. Behav. 61, 299–315.
Dekel, , Fudenberg, E.D., Morris, S., 2006. Topologies on types. Theor. Econ. 1, 275–309.
Dekel, , Fudenberg, E.D., Morris, S., 2007. Interim correlated rationalizability. Theor. Econ. 2, 15–40.
Dufwenberg, M., Stegeman, M., 2002. Existence and uniqueness of maximal reductions under iterated strict dominance. Econometrica 70, 2007–2023.
Ely, J., Peski, M., 2006. Hierarchies of belief and interim rationalizability. Theor. Econ. 1, 19–65.
Fudenberg, D., Levine, D., 1983. Subgame-perfect equilibria of finite and infinite horizon games. J. Econom. Theory 31, 251–268.
Lipman, B., 1994. A note on the implications of common knowledge of rationality. Games Econom. Behav. 6, 114–129.
Mertens, J., Zamir, S., 1985. Formulation of Bayesian analysis for games with incomplete information. Internat. J. Game Theory 10, 619–632.
Osborne, M.J., Rubinstein, A., 1994. A Course in Game Theory. MIT Press.
Pearce, D., 1984. Rationalizable strategic behavior and the problem of perfection. Econometrica 52, 1029–1050.
Tan, T., Werlang, S., 1988. The Bayesian foundations of solution concepts of games. J. Econom. Theory 45, 370–391.
Weinstein, J., Yildiz, M., 2007. A structure theorem for rationalizability with application to robust predictions of refinements. Econometrica 75, 365–400.
Weinstein, J., Yildiz, M., 2013. Robust predictions in infinite-horizon games—An unrefinable folk theorem. Rev. Econom. Stud. 80, 365–394.
Yildiz, M., 2015. Invariance to representation of information. Games Econom. Behav. 94, 142–156.