Uncertain Rationality and Robustness in Games with Incomplete Information

Fabrizio Germano
Peio Zuazo-Garin

March 2015

Barcelona GSE Working Paper Series

Working Paper nº 814

Uncertain Rationality and Robustness in Games with Incomplete Information*

Fabrizio Germano[†] Peio Zuazo-Garin[‡]

February 26, 2015

Abstract

Economic predictions are highly sensitive to model and informational specifications. Weinstein and Yildiz (2007) show that, in static games with incomplete information, only very weak predictions, namely, the interim correlated rationalizable (ICR) actions, are robust to higher-order belief misspecifications. This paper extends their robustness analysis to allow for higher-order uncertainty about rationality. We introduce *interim correlated p-rationalizability* (ICR^p) as a solution concept for games with incomplete information. We first confirm the robustness of the ICR predictions to small departures from common belief in rationality by showing the continuity of ICR^p actions at $p = 1$, where they coincide with ICR. At the same time, we show that Weinstein and Yildiz's (2007) deeper results on the structure of rationalizability, most notably, their discontinuity and generic local uniqueness properties, fail as soon as any arbitrarily small amount of higher-order uncertainty about rationality is introduced. Thus, we find that common belief in rationality is a necessary condition for Weinstein and Yildiz's (2007) discontinuity property to hold. Among other things, this reveals the diminishing strategic impact of higher-order belief constraints.

Keywords: Robustness, rationalizability, uncertain rationality, incomplete information, belief hierarchies.

JEL Classification: C72, D82, D83.

---

*The authors thank Pierpaolo Battigalli, Eddie Dekel, Aviad Heifetz, Tai-Wei Hu, Willemien Kets and Stephen Morris for suggestions and comments, and participants at the 11th Conference on Logic and Foundations of Logic and Game Theory (LOFT 11) for valuable discussion. Germano acknowledges financial support from the Spanish Ministry of Science and Technology (grant ECO2011-28965), as well as from the Barcelona GSE Research Network and the Generalitat de Catalunya. Zuazo-Garin acknowledges financial help from the Spanish Ministry of Economy and Business Support and from the Department of Language Policy and Culture of the Basque Government (grants ECO2012-31346 and IT568-13, respectively), and thanks the WCAS Department of Economics at Northwestern University for generous hospitality.

†Universitat Pompeu Fabra, Departament d'Economia i Empresa, and Barcelona GSE, Ramon Trias Fargas 25–27, E–8005, Barcelona, Spain; fabrizio.germano@upf.edu.

‡Universitat Rovira i Virgili, CREIP and *BRiDGE* Group; Department of Economics, Avinguda de la Universitat 1, 43204, Reus, Spain; peio.zuazo@urv.cat.

1 Introduction

Whether in experiments or in the field, economic behavior frequently exhibits choices that are inconsistent with rationality or common belief in rationality. Moreover, inconsistencies are also often anticipated by economic agents. In this paper, we relax the crucial assumption of common belief in rationality and, instead, study behavior that obtains when there is only *approximate* common belief in rationality, which we represent *via* common p-belief in rationality. We characterize behavior in this case, for any $p \in [0,1]$, and confirm the robustness of the ICR actions to higher order uncertainty about rationality by showing continuity of predicted behavior at $p = 1$. At the same time, we show that Weinstein and Yildiz's (2007) result that whenever a type has multiple rationalizable actions, any of them can be uniquely rationalized, fails as soon as there is an arbitrarily small departure from common belief in rationality, that is, under common p-belief in rationality, for any $p < 1$. Thus, common belief in rationality is revealed as a *necessary condition* for the results on the structure of rationalizability of Weinstein and Yildiz (2007) and their generalization by Penta (2013). We first illustrate the role of common belief in rationality within a simple example before turning to a more detailed discussion of our set-up and contribution.

1.1 Perturbations of common belief in rationality: an example

Assumptions on higher-order beliefs in opponents' rationality can have a significant impact on economic predictions, and in particular, expected behavior can be drastically influenced by arbitrarily small departures from common belief in rationality. Think of Alexei Ivanovich and Polina Alexandrovna, who assume the roles of row and column player, respectively, in the following parametrized 2×2 strategic-form game:

	L	R
T	2, 2	$\frac{2-\theta}{1-\theta}$, 0
B	1, $\frac{-1}{1-\theta}$	$\frac{1}{1-\theta}$, 1

where parameter θ, drawn from $\Theta = [0,1)$, represents some uncertain state of nature. In addition, this game with incomplete information is endowed with the following belief structure: each player believes that ℓ, a lower bound for θ, is commonly known, and when so, each player believes that it is common belief that θ is uniformly distributed on the interval $(\ell, 1)$, for some $\ell \in (0,1)$.[1] Thus, Alexei and Polina's interim beliefs can be identified with the lower bounds ℓ_A and ℓ_P, which each of them believes to be commonly known.

[1]As a brief advance of the notation that will follow in the paper, formally we have sets of types $T_A = T_P = \Theta$, and belief maps given by $\pi_A(t_A)[(t_P,\theta)] = 1_{\{t_A\}}(t_P) \cdot \left(\frac{1}{1-t_A}\right)$, and $\pi_P(t_P)[(t_A,\theta)] = \pi_A(t_P)[(t_A,\theta)]$ for any $(t_A, t_P, \theta) \in T_A \times T_P \times \Theta$.

It is obvious that action T is strictly dominant for Alexei regardless of the specific lower bound ℓ_A, he believes commonly known. Therefore, since L is the only best response to any belief of Polina's that assigns probability 1 to Alexei being rational (no matter which lower bound ℓ_P, Polina believes commonly known), L turns out to be the only interim correlated rationalizable (ICR, Dekel *et al.*, 2007) action for Polina, regardless again of her beliefs in the state of nature. Note that, Polina assigning probability 1 to Alexei being rational, plays a crucial role in this conclusion: to see it, let's assume alternatively that Polina believes in Alexei's rationality with some probability p strictly lower than 1, but still, possibly as close to 1 as desired. In particular, any belief, say, η that is consistent with Polina's interim higher-order beliefs on Θ and assigns probability $(1-p)$ to Alexei choosing B, is also consistent with this assumption. It is easy to check that under such belief η,[2] choosing L yields Polina a lower expected payoff than choosing R, regardless of the lower bound ℓ_P Polina believes to be commonly known, and thus, choice R is rationalized by η. This stands in glaring contrast with the previous case where Polina assigns probability 1 to Alexei being rational, and in which action R happened to be non-rationalizable.

The example above illustrates how model misspecifications involving scenarios with arbitrarily small departures from the assumption, that players commonly believe in each others' rationality, might substantially distort economic predictions. In this particular example, such distortions are indeed independent of players' interim beliefs regarding the state of nature, and, more importantly, arise regardless of how close we are from common belief in rationality (in terms of the distance between p and 1). Thus, the example highlights the importance of investigating when behavior prescribed by interim correlated rationalizability exhibits robustness to small perturbations in higher-order beliefs in rationality, so that model misspecifications of this kind do not drastically influence predictions on players' behavior.[3] This is precisely what the analysis of the present paper contributes to.

1.2 Rationalizability, robustness and Weinstein and Yildiz's discontinuity

The introduction of rationalizability (Bernheim and Pearce, both in 1984) stands as a central contribution to non-cooperative game theory. The use of possibly deluded conjectures allowed for departure from a long standing focus on equilibrium, and the explicit role of beliefs in its formulation contributed to a deeper insight into the underlying epistemic requirements: Brandenburger and Dekel (1987) and Tan and Werlang (1988),[4]

[2] As easy as advancing to paragraph 3.2.2, where a full formal exposition is presented.

[3] As a curiosity, it should be pointed out that indeed, a rational Alexei would *always* (as long as the lower bound he believed to be commonly known was greater than 0) be interested in, if possible, making Polina not believe he is rational. This would be an interesting objective for Alexei in any situation of pre-play communication.

[4] With caveats on allowing for correlated beliefs and the formalism employed.

show that rationalizability characterizes behavior under common belief in rationality in strategic-form games with complete information, and thus, stands as an outpost in the border next to the realm of bounded rationality, and as a relevant benchmark for solution concepts defined for this class of games. Extending rationalizability to incomplete information regarding payoffs, entails numerous additional subtleties. A first approach by Battigalli and Siniscalchi (2003) builds on the novel notion of Δ-*restrictions*, which formalize constraints on players' first order beliefs (on payoff states), to introduce Δ-*rationalizability* as such a generalization.[5] However, it is not immediately clear how Δ-restrictions should be implemented in Bayesian games (originally introduced by Harsanyi, 1967–1968, and the main incarnation of games with incomplete information in applied models of Microeconomic Theory), that is, games with incomplete information with an exogenously fixed type structure that encodes the set of possible belief hierarchies players can hold at interim level. This problem is tackled first by Ely and Pęski (2006) and later, by Dekel *et al.* (2007) who propose *interim independent rationalizability* (IIR) and *interim correlated rationalizability* (ICR), respectively, as solution concepts that specify behavior under common belief in rationality at interim level, that is, once a type has been fixed. Methodological concerns arise, though: by its definition, IIR neglects correlation between payoff states and types, and thus, turns out to be sensitive to *type-redundancy*, that is, *non-robust* to which specific type structure is employed to model some fixed belief hierarchies. Dekel *et al.* (2007) prove that this is not the case with ICR: different types that induce the same hierarchies do have equal sets of ICR actions, and hence, ICR overcomes the problem of type (but not belief hierarchy) misspecification.[6] Finally, recent work by Battigalli *et al.* (2011) presents a full epistemic analysis of rationalizability and incomplete information, performed under canonical type spaces,[7] that establishes the precise relation between Δ-rationalizability and the two notions of interim rationalizability mentioned above.

Apart from the robustness property already mentioned (regarding specific type-representation of belief hierarchies), it is natural to wonder whether ICR is robust with respect to the following two, more traditional, interpretations:

(*i*) Does unexpected behavior arise under arbitrarily small misspecifications of belief hierarchies? Robustness of ICR in this sense is ensured by the negative answer to the last question provided by Dekel *et al.* (2007), who prove that as a correspondence from belief hierarchies to actions, ICR is upper-hemicontinuous.[8] More surprisingly though, it is shown by Weinstein and Yildiz (2007) that ICR is indeed the *sharpest* set of robust predictions when all common knowledge assumptions

[5] Proving by the way, that Δ-rationalizability characterizes common belief in rationality and the Δ-restrictions. It should be noted that their notion of Δ-rationalizability deals with the more general domain of dynamic games, which include normal-form ones as particular case.

[6] It should be added that indeed, belief hierarchies are elicitable, whereas type structures are not.

[7] Or *universal* type space; see Mertens and Zamir (1985) or Brandenburger and Dekel (1993).

[8] This analysis complements the work by Monderer and Samet (1989) and Kajii and Morris (1997), who study robustness of equilibria in setting of complete information modeling higher-order uncertainty about payoffs applying p-belief techniques.

regarding payoffs are relaxed;[9] that is, every refinement of ICR (every notion of equilibria in particular) fails to be robust to misspecifications of belief hierarchies. Specifically, Weinstein and Yildiz (2007) prove that if the set of payoff states is *rich* enough so that for any action there is some state in which the action is strictly dominant (*richness* condition), then, any ICR action is uniquely selected in along some sequence of models converging to the original model. Together with upper-hemicontinuity of ICR, this result implies that the set of belief hierarchies with unique ICR action is open and dense, that is, that ICR generically (in topological terms) selects a unique outcome. Weinstein and Yildiz's (2007) message is reinforced by Penta (2013), who obtains an analogous result in less demanding environments that do allow for some common knowledge assumptions to be maintained.[10]

(*ii*) Are ICR predictions affected by higher-order uncertainty about rationality? The closest version to the analysis of this robustness issue can be found in Hu's (2007) work on compact games with complete information, where common belief in rationality is approximated by common p-belief in rationality, and it is shown that p-*rationalizability*, a solution concept introduced in the paper and proved to characterize common p-belief in rationality, converges to rationalizability as p approximates 1 (*i.e.*, as higher-order in rationality is vanished).

Still, the effects of higher-order uncertainty about rationality, and more importantly, of higher-order joint uncertainty about both rationality and payoff states in games with incomplete information and ICR are not clear. The last issue gains particular relevance when noticing that a key feature of the results by Weinstein and Yildiz (2007) and Penta (2013) is the fact that common belief in rationality enables cascade effects provoked by high-order restrictions of arbitrary order on the belief hierarchy, whose effect is diminishing as higher-order uncertainty about rationality is introduced. Thus, the aforementioned problem regarding the robustness of ICR to higher-order uncertainty about rationality remains open, and an additional one arises naturally: is Weinstein and Yildiz's (2007) discontinuity robust to higher-order uncertainty about rationality?

This paper extends Hu's (2007) formalism to games with incomplete information in order to provide an answer to the two robustness problems mentioned above.[11] Common belief in rationality is approximated *via* common p-belief in rationality and *interim*

[9] *Relaxed* is probably not the most accurate term at this point: notice that the less players commonly know, the more players are somehow commonly aware of; thus, it is not so clear whether we *relax* or *strengthen* our epistemic assumptions.

[10] Penta (2012) and Chen (2012) extend the study of generic uniqueness from normal-form games to extensive-form ones.

[11] Within a partitional framework à la Aumann, Germano and Zuazo-Garin (2015) relax rationality and common knowledge of rationality in complete and incomplete information games and replace them with a substantially weaker notion of common knowledge of mutual p-belief in rationality, which they use to show robustness of the correlated equilibrium and the Bayes correlated equilibrium (Bergemann and Morris, 2013) solution concepts to bounded rationality.

correlated p-rationalizability (ICR^p) is proposed as the solution concept for games with incomplete information that captures rational behavior under the latter epistemic constraint.

We first build on the framework for epistemic analysis due to Battigalli *et al.* (2011) to provide an epistemic foundation of ICR^p, namely, that it characterizes rationality and common p-belief in rationality (as originally intended) at interim level (Theorem 1). This epistemic characterization legitimates ICR^p as a suitable theoretical tool for the formalization of behavioral implication due to departures from common belief in rationality, modeled by common p-belief in rationality, and thus, enables us to proceed to the study of our robustness concerns. Next we check three elementary robustness properties. First, we present an alternative definition of ICR^p for Bayesian games in terms if types, and show that the set of ICR^p obtained for a type under this second definition exactly coincides with the set of ICR^p actions obtained for the belief hierarchy induced by the type under the original definition (Proposition 1). Thus, we conclude that ICR^p is robust to specific type-representation of belief hierarchies (and in particular, that so is ICR, which coincides with ICR^p for $p = 1$). Second, we check upper-hemicontinuity on belief hierarchies of our solution concept (Proposition 2; of course, upper-hemicontinuity of ICR follows). Finally, we obtain the robustness of ICR to higher-order uncertainty about rationality by proving in Proposition 3 that ICR^p is an upper-hemicontinuous correspondence on p for each fixed belief hierarchy, and continuous indeed at $p = 1$.

We then go on to the main result of the paper, which concerns a key structural property of Weinstein and Yildiz (2007). We prove in Theorem 2 that, given any belief hierarchy, it holds that, for any $p < 1$, the set of interim correlated *strictly* rationalizable actions (ICSR; the analogous of ICR in which elimination surviving actions are required to be the unique best response to some suitable conjecture) corresponding to the hierarchy is included in the set of ICR^p actions corresponding to every belief hierarchy in some tail of every sequence converging to the original hierarchy. Despite the rather technical nature of the result, its implications for Weinstein and Yildiz's (2007) discontinuity are immediate: in case there exists some belief hierarchy with multiple ICSR,[12] then, no sequence converging to this belief hierarchy admits a unique ICR^p selection all along the sequence, and since this impossibility of asymptotic unique selection holds for any p arbitrarily close (but different) from 1, it follows that Weinstein and Yildiz's (2007) key discontinuity property is non-robust to higher-order uncertainty about rationality. The non-robustness of unique selection results from the global games literature also follows. Accordingly, these results highlight the diminishing impact of higher-order belief restrictions as higher-order uncertainty about rationality is introduced; this is natural if we notice that, under p-belief in rationality, we assume that opponents will play their (possibly) unique rational choice with at least probability p, but the probability we assign to them playing some best response to rational choices of ours is already p^2, which is strictly lower than p. This is reminiscent of the analysis by Chen *et al.* (2010), who unlike the present paper, abandon the product topology and propose an alternative one

[12] Obviously, only trivial games with incomplete information do not admit such belief hierarchy.

(the *uniform-weak topology*) which focuses on similarity of lower-order beliefs, and thus accounts for this effect of the p-belief approach.

Finally, we present some examples that illustrate the robustness issues studied in abstract terms throughout the paper. Among these, we show that Rubinstein's (1989) *Email game* is especially robust to higher-order uncertainty: players' behavior remains as predicted by the unique ICR action profile for any p above $1/2$.

The rest of the paper is structured as follows: Section 2 recalls the formalism, standard in game-theoretical literature, needed for analysis, which consists of games with incomplete information, construction of the canonical (or universal) type space, and a brief review of rationalizability in games with incomplete information, and of Weinstein and Yildiz's discontinuity. Section 3 presents our approximation to rationalizability in games with incomplete information, the non-robustness of Weinstein and Yildiz (2007)'s discontinuity result, the result concerning robustness of ICR, and three examples. Section 4 ends with some additional properties of approximate rationalizability and some comments.

2 Preliminaries

2.1 Games with incomplete information

A (*static*) *game with incomplete information* consists of a list $\mathscr{G} = \left\langle I, \Theta, (A_i, u_i)_{i\in I} \right\rangle$, where: (*i*) I is a *finite* set of players, (*ii*) Θ is a compact and metrizable set of *payoff states*, and for each player i we have (*iii*) a compact and metrizable set of *actions*, A_i, and (*iv*) a continuous *utility map* $u_i : A \times \Theta \to \mathbb{R}$, where $A = \prod_{i\in I} A_i$ is the set of action *profiles*. For each player i, we refer to a probability measure $\mu \in \Delta(A_{-i} \times \Theta)$, where $A_{-i} = \prod_{j\neq i} A_j$, as a *conjecture*,[13] and we define player i's *best reply correspondence* as,

$$\begin{array}{rccc} BR_i: & \Delta(A_{-i}\times\Theta) & \rightrightarrows & A_i \\ & \mu & \mapsto & \underset{a_i\in A_i}{\operatorname{argmax}} \int_{A_{-i}\times\Theta} u_i\left((a_{-i};a_i),\theta\right) d\mu, \end{array}$$

which, due to the topological assumptions specified above, is known to be both non-empty and upper-hemicontinuous.[14] Following Harsanyi's (1967–1968) approach, it is usual in economic literature to model agents' probabilistic private information over the set of payoff states by explicitly imposing implicit restrictions on higher-order beliefs

[13] For a given topological space X we denote by $\Delta(X)$ the space of all probability measures on the Borel subsets of X endowed with the weak$*$ topology, so that if X is compact and metrizable, so is $\Delta(X)$. In particular, every continuous function under this topology will be measurable under the corresponding Borel σ-algebra, $\mathscr{B}(X)$. Topologies for other kind of spaces are standard: the induced topology for subsets and the Tychonoff topology for products.

[14] When necessary, with some abuse of notation we will write $BR_i(\mu) = BR_i\left(\operatorname{marg}_{A_{-i}\times\Theta}\mu\right)$ for any belief $\mu \in \Delta(X \times A_{-i} \times \Theta)$, being X some compact and metrizable space.

on Θ.[15] This is achieved by introducing a *standard type structure*, which consists of a list $\mathscr{T} = \langle T_i, \pi_i \rangle_{i \in I}$, where for each player i we have: (i) a compact and metrizable set of *standard types*, T_i, and (ii) a continuous *belief map* $\pi_i : T_i \to \Delta(T_{-i} \times \Theta)$, where $T_{-i} = \prod_{j \neq i} T_j$. We refer to a pair $\Gamma = \langle \mathscr{G}, \mathscr{T} \rangle$ as a *Bayesian game.*

2.2 Canonical types

Despite its popularity in economic literature, standard types have some drawbacks for epistemic analysis; in particular: (i) they are defined as elements of an exogenous *supra*-structure added to the game, rather than in terms of the components of the game themselves, and (ii) type structures are not elicitable. For this reason, we adhere to the approach by Battigalli *et al.* (2011) and Dekel and Siniscalchi (2013),[16] based on belief hierarchies, and recall the construction of canonical type spaces,[17] upon which the our analysis and results are formulated.

2.2.1 Construction

We follow Brandenburger and Dekel's (1993) construction of canonical type spaces. For each player i, fix (compact and metrizable) basic uncertainty space U_i. Then, let $X_i^1 = U_i$, and set recursively, $X_i^n = X_i^{n-1} \times \prod_{j \neq i} \Delta\left(X_j^{n-1}\right)$ for each $n \geq 2$. We refer to each element $h_i \in \mathcal{H}_i^0 = \prod_{n \in \mathbb{N}} \Delta(X_i^n)$ as a *belief hierarchy*, and we refer to the n^{th} component of h_i as the n^{th} *order belief* of player i. We say that belief hierarchy h_i is *coherent* if higher-order belief do not contradict lower order ones, *i.e.*, if $\text{marg}_{X_i^n} h_i^{n+1} = h_i^n$ for any $n \in \mathbb{N}$, and we denote the set of coherent belief hierarchies by $\mathcal{H}_i^1$. Now, let,[18]

$$\mathcal{H}_i^{n+1} = \left\{ h_i \in \mathcal{H}_i^n \,\middle|\, h_i^m \left[\text{Proj}_{X_i^m}\left(\mathcal{H}_{-i}^n \times U_i\right)\right] = 1 \text{ for any } m \in \mathbb{N} \right\},$$

where $\mathcal{H}_{-i}^n = \prod_{j \neq i} \mathcal{H}_j^n$, is recursively defined for any $n \in \mathbb{N}$, and set $\mathcal{H}_i = \bigcap_{n \geq 0} \mathcal{H}_i^n$. Belief hierarchies in $\mathcal{H}_i$ can be understood as those satisfying common belief in coherency,[19] and they are the ones to which we will refer to exclusively, with some abuse of language, when we invoke the term "belief hierarchy." In addition, Brandenburger and Dekel (1993) prove that there exists a homeomorphism $\upsilon_i : \mathcal{H}_i \to \Delta(\mathcal{H}_{-i} \times U_i)$, where $\mathcal{H}_{-i} = \prod_{j \neq i} \mathcal{H}_j$, such that $\text{marg}_{X_i^n} \upsilon_i(h_i) = h_i^n$ for any belief hierarchy h_i, and any $n \in \mathbb{N}$. Thus, in terms of *expressible* events, we are enabled to work on either space,

[15] A brief discussion on the (possibly controversial) use of the expression *private information* can be found on Section 4, paragraph C.

[16] Who the reader is referred to for accurate discussions on critiques (i) and (ii), respectively.

[17] Most commonly –*i.e.*, frequently– known as *universal type spaces* in the literature; see Armbruster and Böge (1979), Böge and Eisele (1979), Mertens and Zamir (1985) or Brandenburger and Dekel (1993) for the case with topological assumptions, or Heifetz and Samet (1998) for the purely measure-theoretic one.

[18] For any product space $X \times Y$ and any subset $S \subseteq X \times Y$, we denote *projections* on some component of X by $\text{Proj}_X S = \{x \in X \,|\, (x, y) \in S \text{ for some } y \in Y\}$.

[19] Epistemic notions such as *common belief* are properly formalized in Section 2.3.

$\mathcal{H}_i$ or $\Delta(\mathcal{H}_{-i} \times U_i)$, depending on which is more convenient each time. We refer to $(\mathcal{H}_i, \upsilon_i)_{i \in I}$ as the *canonical type spaces* on basic uncertainty spaces $(U_i)_{i \in I}$. The analysis of payoff uncertainty focuses on canonical type space $(\mathcal{T}_i, \varphi_i)_{i \in I}$ with basic uncertainty space $U_i = \Theta$ for any player i, and we refer to each canonical type $\tau_i \in \mathcal{T}_i$ as a *standard hierarchy*.

2.2.2 Canonicity and standardness

The fact that, as mentioned above, we opted to work with standard hierarchies rather than standard types begs for a clarification of the relation between the two. It is known that every standard type of any fixed standard type structure induces a unique standard hierarchy, not necessarily in an injective way, though.[20] Such induced standard hierarchies are obtained in the following way: for each standard type t_i let $\tau_{i,0}(t_i) = \text{marg}_\Theta \pi_i(t_i)$, and for each $n \in \mathbb{N}$ and $E_n \in \mathscr{B}\left(\text{Proj}_{X^n_{-i}}(\mathcal{T}_{-i} \times \Theta)\right)$, let:

$$\tau_{i,n}(t_i)[E_n] = \pi_i(t_i)\left[\{(t_{-i}, \theta) \in T_{-i} \times \Theta \,|(\tau_{-i,n-1}(t_{-i}), \theta) \in E_n\}\right].$$

Note that it is not clear whether this definition is right, in the sense that it might be the case that some of the sets we consider on the right side are not measurable; this is fortunately not the case, and it is known from Mertens and Zamir (1985) or Brandenburger and Dekel (1993) among others, that map τ_i is not only well-defined, but also continuous for any player i, and that in addition, the standard hierarchies induced do indeed represent common belief in coherency.

2.3 Epistemic framework

If we apply Brandenburger and Dekel's (1993) construction to the family of basic uncertainty spaces $(A_{-i} \times \Theta)_{i \in I}$, we obtain canonical type space $(\mathcal{E}_i, \psi_i)_{i \in I}$, whose corresponding belief hierarchies we refer to as *epistemic hierarchies*. This way, following Battigalli *et al.* (2011), the epistemic analysis is based on epistemic hierarchies, and performed in state space $\Omega = \mathcal{E} \times A \times \Theta$, where $\mathcal{E} = \prod_{i \in I} \mathcal{E}_i$. For each player i we denote $\Omega_i = \mathcal{E}_i \times A_i$, and for each state ω, we will consider the following projections: $\omega_i = \text{Proj}_{\Omega_i}(\omega)$, $e_i(\omega) = \text{Proj}_{\mathcal{E}_i}(\omega)$, $a_i(\omega) = \text{Proj}_{A_i}(\omega)$ and $\theta(\omega) = \text{Proj}_\Theta(\omega)$. Thus, each state is a description of payoff states, players' choices, and players' belief hierarchies on the previous two contingencies. The epistemic language is completed as follows.

2.3.1 Rationality and common p-belief in rationality

We say that payer i is rational at state ω whenever her choice at ω is optimal given her beliefs at ω. Thus, $R_i = \{\omega \in \Omega \,|a_i(\omega) \in BR_i(e_{i,1}(\omega))\}$ represents the event that *player i is rational*. As usual, we denote $R = \bigcap_{i \in I} R_i$ and $R_{-i} = \bigcap_{i \in I} R_i$. Note that each R_i is closed due to BR_i being closed-valued, and Proj_{A_i}, continuous. Following Monderer

[20] Giving rise the *type redundancy*, when injectiveness is not satisfied.

and Samet (1989), assumptions on players' beliefs are represented employing p-belief operators. Formally, for positive probability p, player i's *p-belief operator* is defined as map $E \mapsto B_i^p(E)$, where for any event E,

$$B_i^p(E) = \left\{ \omega \in \Omega \,\middle|\, \psi_i(e_i(\omega)) \left[(\omega'_{-i}, \theta) \in \mathcal{E}_{-i} \times A_{-i} \times \Theta \,\middle|\, (\omega'_{-i}, \omega_i, \theta) \in E \right] = 1 \right\}.$$

That is, event $B_i^p(E)$ is the collection of states in which player i assigns at least probability p to event E, and we refer to it as the event that *player i p-believes E*. The *mutual p-belief operator* is given by $E \mapsto B^p(E) = \bigcap_{i \in I} B_i^p(E)$, for any event E. When p equals 1 we drop superscripts and refer to 1-belief as simply, *belief*. Note that from the fact that every ψ_i is a homeomorphism, it follows that p-belief operators are closed-valued. Finally, higher-order epistemic restrictions are imposed *via* the *common p-belief operator*, which is recursively defined as follows: let $CB^p(E) = \bigcap_{n \in \mathbb{N}} B^{n,p}(E)$, where we set $B^{1,p}(E) = B^p(E)$, and for any $n \in \mathbb{N}$, $B^{n+1,p}(E) = B^p(B^{n,p}(E))$. Again, we write $CB(E) = CB^1(E)$ for any event E to represent *common belief*. Following the approach by Monderer and Samet (1989), Kajii and Morris (1997) or Hu (2007) among others, we later employ common p-belief to model arbitrarily small departures form common belief.

2.3.2 From epistemic to standard hierarchies

As shown by Battigalli *et al.* (2011), it is possible to construct, by recursive marginalization, quotient maps $q_i : \mathcal{E}_i \to \mathcal{T}_i$ and $\bar{q}_i : \Delta(\mathcal{E}_{-i} \times A_{-i} \times \Theta) \to \Delta(\mathcal{T}_{-i} \times \Theta)$ that make the following diagram commutative:

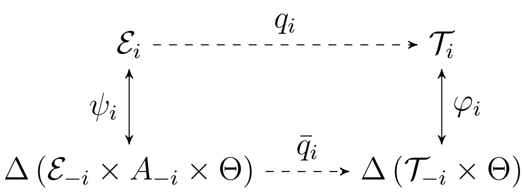

so that consistency between events that are expressible in each domain, the ones corresponding to uncertainty about Θ and uncertainty about $A_{-i} \times \Theta$, is guaranteed. Then, for any player i and standard hierarchy τ_i, we set $[q_i = \tau_i] = \{\omega \in \Omega \,|\, q_i(e_i(\omega)) = \tau_i\}$ as the event that player i's standard hierarchy is exactly τ_i. Note that $[q_i = \tau_i]$ is well-defined and also closed, due to q_i being continuous.

2.4 Rationalizability and Weinstein and Yildiz's discontinuity

The standard notion of interim rationalizability employed in Bayesian games is *interim correlated rationalizability* (ICR), introduced by Dekel *et al.* (2007). Despite their original definition being given in terms of standard types, the authors show (Proposition 1 in their paper) that ICR does not really depend on each standard type, but rather, on the infinite hierarchy of beliefs induced by the standard type;[21] a fact that implicitly suggests that an analogous notion of ICR can be found for games with incomplete information (*i.e.*, in the absence of any specific standard type structure). Such definition is provided in the exhaustive analysis of rationalizability for games with incomplete information by Battigalli *et al.* (2011).[22] Since our results in Section 3 are proved employing standard hierarchies, we adhere to this latter version of ICR. Formally, player i's standard hierarchy τ_i's set of interim correlated rationalizable actions is iteratively defined as follows: let, $\mathrm{ICR}_{i,0}(\tau_i) = A_i$ and $C_{i,0}(\tau_i) = \left\{\eta \in \Delta(\mathcal{T}_{-i} \times A_{-i} \times \Theta) \,\middle|\, \mathrm{marg}_{\mathcal{T}_{-i}\times\Theta}\eta = \varphi_i(\tau_i)\right\}$, and for any $n \in \mathbb{N}$,[23]

$$\mathrm{ICR}_{i,n}(\tau_i) = \{a_i \in A_i \,|a_i \in BR_i(\eta) \text{ for some } \eta \in C_{i,n-1}(\tau_i)\},$$
$$C_{i,n}(\tau_i) = \{\eta \in C_{i,n-1}(\tau_i) \,|\eta[\Theta \times \mathrm{Graph}(\mathrm{ICR}_{-i,n})] = 1\},$$

then, player i's standard type τ_i's set of *interim correlated rationalizable* actions is defined as $\mathrm{ICR}_i(\tau_i) = \bigcap_{n\geq 0} \mathrm{ICR}_{i,n}(\tau_i)$. In addition to type-representation independence, each ICR_i can be proved to be upper-hemicontinuous in τ_i, and as shown by Dekel *et al.* (2007) and Battigalli *et al.* (2011), it characterizes rational behavior under common belief in rationality, that is, $\mathrm{ICR}_i(\tau_i) = \mathrm{Proj}_{A_i}(R_i \cap CB(R) \cap [q_i = \tau_i])$ for any player i and standard hierarchy τ_i.[24]

In their the study of ICR, Weinstein and Yildiz (2007) find the following surprising *discontinuity* property of ICR for games with finite set of action profiles: if Θ is such that there is some strict dominance state for any action of every player (the *richness condition*), then, for any standard hierarchy τ_i, and any $a_i \in \mathrm{ICR}_i(\tau_i)$, there is some sequence $(\tau_i^n)_{n\in\mathbb{N}}$ converging to τ_i such that $\mathrm{ICR}_i^i(\tau_i^n) = \{a_i\}$.[25] We refer to this property as *Weinstein and Yildiz's discontinuity*. In particular, this result, together with the upper-hemicontinuity of each ICR_i implies *generic uniqueness* under the richness condition, *i.e.*, that the set of each player i's standard hierarchies with a unique ICR action is open and dense in $\mathcal{T}_i$.

---

[21]We refer to this property as *type-representation independence*. This is not satisfied by another well-known notion of interim rationalizability: Ely and Pęski's (2006) *interim independent rationalizability* (IIR). The difference between ICR and IIR lies in the fact that the latter imposes that players' beliefs assume that payoff states and opponents' choices are uncorrelated; an assumption not *expressible* in the language of belief hierarchies (see Battigalli *et al.*, 2011).

[22]Who in addition, show how ICR can be regarded as a particular case of Battigalli and Siniscalchi's (2003; 2007) notion of Δ-*rationalizability*.

[23]We introduce some notation below: since for each player i, and each $n \geq 0$, $\mathrm{ICR}_{i,n}$ can be regarded as a correspondence from $\mathcal{T}_i$ to A_i, and we denote $\mathrm{ICR}_{-i,n}(\tau_{-i}) = \prod_{j\neq i} \mathrm{ICR}_{j,n}(\tau_j)$, we abbreviate as follows: $\mathrm{Graph}(\mathrm{ICR}_{-i,n}) = \{(\tau_{-i}, a_{-i}) \in \mathcal{T}_{-i} \times A_{-i} \,|\, a_{-i} \in \mathrm{ICR}_{-i,n}(\tau_{-i})\}$. Note that we can *play* with the idea of, with some lack of accuracy, identifying support and probability 1 due to the fact that we later integrate in order to compute expected payoffs.

[24]Some formal inaccuracy should be noted here: $CB(R)$ already implies R; still, we keep the redundant notation and terminology for expositional emphasis.

[25]Later, Penta (2013) proved that the rather demanding richness condition can be abandoned, and the discontinuity result extended to relatively mild relaxations of common knowledge assumptions.

3 Uncertain of rationality and robustness

In this section we introduce our notion of approximate rationalizability in games with incomplete information, which is in the spirit of Hu's (2007) p-rationalizability for the complete information case. We apply this notion of approximate rationalizability to robustness properties of ICR to higher-order uncertainty about rationality. In particular, we show that ICR is robust to higher-order uncertainty about rationality, and more importantly, that Weinstein and Yildiz's (2007) discontinuity is non-robust to higher-order uncertainty about rationality; *i.e.*, common belief in rationality is a *sine qua non* condition for their result to hold.

3.1 Interim correlated p-rationalizability

3.1.1 Definition

We approximate ICR under small perturbations in common belief in rationality *via* interim correlated p-rationalizability (ICR^p). As mentioned, this concept generalizes Hu's (2007) notion of approximate rationalizability for games with complete information, to the case of incomplete information. The logic behind the definition of ICR^p resembles that corresponding to ICR, with which it exactly coincides when p equals 1, but with the addition that the possibility of higher-order uncertainty about opponents' rationality is introduced.

Definition 1 (Interim correlated p-rationalizability). *Let $\mathscr{G}$ be a game with incomplete information, and p, a positive probability. Set for each player i and hierarchy τ_i, $\mathrm{ICR}^p_{i,0}(\tau_i) = A_i$ and $C^p_{i,0}(\tau_i) = \left\{\eta \in \Delta(\mathcal{T}_{-i} \times A_{-i} \times \Theta) \,\middle|\, \mathrm{marg}_{\mathcal{T}_{-i}\times\Theta}\eta = \varphi_i(\tau_i)\right\}$, and define recursively, for any $n \in \mathbb{N}$,*

$$\mathrm{ICR}^p_{i,n}(\tau_i) = \left\{a_i \in A_i \,\middle|\, a_i \in BR_i(\eta) \text{ for some } \eta \in C^p_{i,n-1}(\tau_i)\right\},$$
$$C^p_{i,n}(\tau_i) = \left\{\eta \in C^p_{i,n-1}(\tau_i) \,\middle|\, \eta\left[\Theta \times \mathrm{Graph}\left(\mathrm{ICR}^p_{-i,n}\right)\right] \geq p\right\},$$

Then, we say that action a_i is interim correlated p-rationalizable (ICR^p) *for standard hierarchy τ_i, if $a_i \in \mathrm{ICR}^p_i(\tau_i) = \bigcap_{n\geq 0} \mathrm{ICR}^p_{i,n}(\tau_i)$.*

3.1.2 Epistemic foundation

First, we want to make sure that ICR^p is a suitable tool for capturing changes in rational behavior due to the introduction of higher-order uncertainty about rationality. This is proved in the following epistemic characterization result:

Theorem 1 (Epistemic foundation of ICR^p). *Let $\mathscr{G}$ be a game with incomplete information, and p, a positive probability. Then, interim correlated p-rationalizability characterizes rationality and common p-belief in rationality; i.e., for any player i and any* *standard hierarchy* τ_i,

$$\mathrm{ICR}_i^p(\tau_i) = \mathrm{Proj}_{A_i}\left(R_i \cap B_i(CB^p(R)) \cap [q_i = \tau_i]\right).$$

In particular, Theorem 1 straightforwardly implies the two following well-known epistemic characterization results, where by $R_i^p(\mathscr{G}_\theta)$ we denote player i's set of *p-rationalizable* actions, as defined by Hu (2007) for the game with complete information $\mathscr{G}_\theta$ corresponding some commonly known payoff state θ.

Corollary 1 (cf. Dekel *et al.* (2007) and Battigalli *et al.* (2011), and Hu (2007), respectively). *Let $\mathcal{G}$ be a game with incomplete information. Then, for any player i we have:*

1. *For any standard hierarchy* τ_i, $\mathrm{ICR}_i(\tau_i) = \mathrm{Proj}_{A_i}\left(R_i \cap CB(R) \cap [q_i = \tau_i]\right)$.
2. *For any payoff state* θ, $R_i^p(\mathscr{G}_\theta) = \mathrm{Proj}_{A_i}\left(R_i \cap CB^p(R) \cap \left[q_i = \tau_i^\theta\right]\right)$.

3.1.3 Elementary robustness properties

Obviously, it holds that when p equals 1, $\mathrm{IRC}_i^p(\tau_i)$=$\mathrm{IRC}_i(\tau_i)$ for any player i and standard hierarchy τ_i. We show next that ICR^p satisfies two well-known robustness properties: (i) it is robust to type-representation; *i.e.*, two different standard types that induce the same standard hierarchies have the same sets of ICR^p actions, and (ii) it is robust to misspecifications on the standard hierarchy, *i.e.*, upper-hemicontinuous. Finally, we check that ICR is robust to higher-order uncertainty about rationality. To study robustness to type-representation, we first need to introduce a definition of ICR^p in terms of standard types:

Definition 2 (Interim correlated p-rationalizability, standard type version). *Let* Γ *be a Bayesian game, and* p*, a positive probability. For each player* i *and standard type* t_i*, let* $\mathrm{ICR}_{i,0}^{p,\mathscr{T}}(t_i) = A_i$ *and* $C_{i,0}^{p,\mathscr{T}}(t_i) = \left\{\mu \in \Delta(T_{-i} \times A_{-i} \times \Theta) \,\middle|\, \mathrm{marg}_{T_{-i}\times\Theta}\mu = \pi_i(t_i)\right\}$*, and define recursively, for any* $n \in \mathbb{N}$,

$$\mathrm{ICR}_{i,n}^{p,\mathscr{T}}(t_i) = \left\{a_i \in A_i \,\middle|\, a_i \in BR_i(\mu) \text{ for some } \mu \in C_{i,n-1}^{p,\mathscr{T}}(t_i)\right\},$$
$$C_{i,n}^{p,\mathscr{T}}(t_i) = \left\{\mu \in C_{i,n-1}^{p,\mathscr{T}}(t_i) \,\middle|\, \mu\left[\Theta \times \mathrm{Graph}\left(\mathrm{ICR}_{-i,n}^{p,\mathscr{T}}\right)\right] \geq p\right\},$$

Then, we say that action a_i *is* interim correlated p-rationalizable (ICR^p) *for standard type* t_i*, if* $a_i \in \mathrm{ICR}_i^{p,\mathscr{T}}(t_i) = \bigcap_{n\geq 0} \mathrm{ICR}_{i,n}^{p,\mathscr{T}}(t_i)$.

Then, it is easy to check that both the definitions of ICR^p are consistent; or in other words, that ICR^p is *type-representation invariant*: the specific type structure representation that encodes the hierarchies is non-relevant:

PROPOSITION 1 (Type-representation independence). *Let* Γ *be a Bayesian game. Then, for any player* i*, any standard type* t_i*, any positive probability* p *and any non-negative integer* n*, it holds that* $\mathrm{ICR}_{i,n}^{p,\mathscr{T}}(t_i) = \mathrm{ICR}_{i,n}^{p}(\tau_i(t_i))$.

In addition, arbitrarily small misspecifications in belief hierarchies do not lead to unexpected radically different behavior:

PROPOSITION 2 (Robustness to uncertainty about payoffs). *Let* $\mathscr{G}$ *be a game with incomplete information, and* p*, a positive probability. Then, for any player* i*, correspondence* $\mathrm{ICR}_i^p : \mathcal{T}_i \rightrightarrows A_i$ *is upper-hemicontinuous.*

We finish the analysis of elementary robustness properties extending Hu's (2007) analysis of robustness of rationalizability to higher-order uncertainty about rationality from games with complete information, to the incomplete information case. Since we already checked in Theorem 1 that ICR^p captures changes in rational behavior corresponding to the introduction of higher-order uncertainty about rationality, the following continuity result in Proposition 3 lets us conclude that ICR is indeed robust to higher-order uncertainty about rationality.[26]

PROPOSITION 3 (Continuity properties on p). *Let* $\mathscr{G}$ *be a game with incomplete information. Then, for any player* i *and any standard hierarchy* τ_i*, correspondence given by* $p \mapsto \mathrm{ICR}_i^p(\tau_i)$ *is upper-hemicontinuous, continuous at* $p = 1$*, and satisfies that* $1 \mapsto \mathrm{ICR}_i(\tau_i)$.

3.2 Non-robustness of Weinstein and Yildiz's discontinuity

We present now the main result of the paper. Remember, as mentioned in Section 2.4, that Weinstein and Yildiz (2007) prove in their Proposition 1, that given a game with incomplete information with finite sets of actions $\mathscr{G}$, if *richness* is satisfied,[27] then, for any player i, any standard hierarchy $\tau_i \in \mathcal{T}_i$ and any $a_i \in \mathrm{ICR}_i(\tau_i)$, then, there exists a convergent sequence $(\tau_i^n)_{n\in\mathbb{N}}$ with limit τ_i, and such that $\mathrm{ICR}_i(\tau_i^n) = \{a_i\}$ for any $n \in \mathbb{N}$. In particular, this structure theorem for rationalizability implies *generic uniqueness* of rationalizability, *i.e.*, that the following set is open and dense in $\mathcal{T}_i$:

$$\mathcal{U}_i = \{\tau_i \in \mathcal{T}_i \,||\mathrm{ICR}_i(\tau_i)| = 1\}.$$

In order to formulate our result, we need to recall the following auxiliary notion both Weinstein and Yildiz (2007) and Penta (2013) make use of in the proofs of their main results.

[26] A related result by Germano and Zuazo-Garin (2015) shows that their notion of *p-rational outcomes* (which coincide with the correlated equilibria when $p = 1$) are continuous in p, for any $p \leq 1$, which, in particular, implies robustness of correlated equilibria to bounded rationality.

[27] That is, for any player i and any action A_i, there is some $\theta \in \Theta$ in which a_i is strictly dominant. Penta (2013) presents a substantial weakening of this assumption.

Player i's standard hierarchy τ_i's set of interim correlated rationalizable actions is recursively defined as follows; let, $\mathrm{ICSR}_{i,0}(\tau_i) = A_i$ and $D_{i,0}(\tau_i) = \{\eta \in \Delta(\mathcal{T}_{-i} \times A_{-i} \times \Theta) \mid \mathrm{marg}_{\mathcal{T}_{-i}\times\Theta}\eta = \varphi_i(\tau_i)\}$, and for any $n \in \mathbb{N}$,

$$\mathrm{ICSR}_{i,n}(\tau_i) = \{a_i \in A_i \,|BR_i(\eta) = \{a_i\} \text{ for some } \eta \in D_{i,n-1}(\tau_i)\},$$
$$D_{i,n}(\tau_i) = \{\eta \in D_{i,n-1}(\tau_i) \,|\eta[\Theta \times \mathrm{Graph}(\mathrm{ICSR}_{-i,n})] = 1\},$$

then, player i's standard type τ_i's set of *interim correlated strictly rationalizable* actions is defined as $\mathrm{ICSR}_i(\tau_i) = \bigcap_{n\geq 0} \mathrm{ICSR}_{i,n}(\tau_i)$. Note that it might perfectly be the case that the set of ICSR actions is empty for some standard hierarchy τ_i, either because no action is a strict best reply to some belief (a rather pathological case), or because $\varphi_i(\tau_i)$ does not assign probability 1 to a set of opponents' standard hierarchies with non-empty set of ICSR actions. We are now ready to state the following result that shows how Weinstein and Yildiz's (2007) discontinuity for games with finite set of action profiles is non-robust to arbitrarily small perturbations in common belief in rationality:

Theorem 2 (Non-robustness of Weinstein and Yildiz's discontinuity). *Let $\mathscr{G}$ be a game with incomplete information and finite set of action profiles. Then for any standard hierarchy τ_i, any $p < 1$ and any sequence $(\tau_i^n)_{n\in\mathbb{N}}$ converging to τ_i, there is some $n_p \in \mathbb{N}$ such that $\mathrm{ICSR}_i(\tau_i) \subseteq \mathrm{ICR}_i^p(\tau_i^m)$ for any $m \geq n_p$.*

That is, as long as we can find some standard hierarchy τ_i such that $|\mathrm{ICSR}(\tau_i)| \geq 2$,[28] Weinstein and Yildiz's discontinuity property fails to be obtained at τ_i for any $p < 1$, or, in other words, *any arbitrarily small departure from common belief in rationality implies that a Weinstein and Yildiz selection is impossible to perform for standard hierarchies that have more than one rationalizable action.* This readily implies the non-robustness of the global games selection techniques in the following sense: whenever the original game has multiple equilibria, perturbations of belief hierarchies that relax common knowledge assumptions about the payoffs, if coupled with arbitrarily small departures from common belief in rationality, cannot allow for unique selection. Once again, common belief in rationality is necessary to obtain a unique selection. Notice also that unlike the results by Weinstein and Yildiz (2007) and Penta (2013), Theorem 2 does not rely on any richness assumption on the set of payoff-states. Thus, in particular, for any $p < 1$, and even under the richness assumption, it is possible to find refinements of every ICR^p that are robust. ICR is an obvious one.[29] Thus, an interesting question arises: for any given $p < 1$, is it possible to find some solution concept that: (i) is robust, (ii) refines ICR^p and (iii) is not some ICR^q with $q > p$? Finally, note that Theorem 2 implies that generic uniqueness is not satisfied in non-trivial cases:

Corollary 2 (Non-robustness of generic uniqueness). *Let $\mathscr{G}$ be a game with incomplete information and finite set of action profiles. Then, for any player i for which there exists*

[28] If the game has two strict Bayesian Nash equilibria, for instance.
[29] Or more generally, any ICR^q with $q > p$.

some τ_i such that $|\mathrm{ICSR}_i(\tau_i)| > 1$, the following set is not dense for any $p < 1$,

$$\mathcal{U}_i^p = \{\tau_i \in \mathcal{T}_i \, | |\mathrm{ICR}_i^p(\tau_i)| = 1\}.$$

3.3 EXAMPLES

We end the section presenting some examples that illustrate the robustness results in Theorem 2 and Proposition 3.

3.3.1 Failure of Weinstein and Yildiz's discontinuity

We begin with the result in Theorem 2. Consider the following two player game with $\Theta = \mathbb{R}$, $A_1 = A_2 = \{\alpha, \beta\}$ and payoffs given by:

	α	β
α	θ, θ	$\theta - 1$, 0
β	0, $\theta - 1$	0, 0

Take some payoff state $\bar{\theta} \in (0, 1)$, that is, one in which the game with complete information induced has two strict Nash equilibria, and consider $\tau_1^{\bar{\theta}}$, the standard hierarchy that represents common belief in $\bar{\theta}$ for player 1 (row). Fix some $p < 1$ and consider any sequence $(\tau_1^n)_{n \in \mathbb{N}}$ converging to $\tau_1^{\bar{\theta}}$. It is easy to see that there is some natural N_p such that $\varphi_1(\tau_1^m)\left[\left(\tau_2^{\bar{\theta}}, \bar{\theta}\right)\right] > p$ for any $m \geq N_p$. Now, for any natural n, let belief $\eta^{n,q} \in \Delta(\mathcal{T}_2 \times A_2 \times \Theta)$ be defined as $\eta^{n,q} = \varphi_1(\tau_1^n) \cdot \left[(1-q) \cdot 1_{\{\alpha\}} + q \cdot 1_{\{\beta\}}\right]$ for any $q \in [0, 1]$. It is clear that $\eta^{n,q}$ induces τ_1^n for any $n \in \mathbb{N}$, and that it satisfies $\eta^{m,q}[\Theta \times \mathrm{Graph}(\mathrm{ICR}_2^p)] > p$ for any $m \geq N_p$. Now, it also holds for any natural n that $\alpha \in BR_1(\eta^{n,0})$ and $\beta \in BR_1(\eta^{n,1})$ if and only if,[30]

$$\int_\Theta \theta \mathrm{d}\left(\mathrm{marg}_\Theta \varphi_1(\tau_1^n)\right) \in (0, 1). \tag{1}$$

But the fact that,

$$\int_\Theta \theta \mathrm{d}\left(\mathrm{marg}_\Theta \varphi_1(\tau_1^n)\right) \overset{n\to\infty}{\longrightarrow} \bar{\theta}$$

guarantees that there is some natural N such that (1) holds for any $m \geq N$. Thus, we have that $\mathrm{ICR}_1^p(\tau_1^m) = A_1$ for any $m \geq n_p = \max\{N_p, N\}$.

[30]This follows from the fact that $u_1(\eta^{n,q}, \alpha) = \int_\Theta [(1-q) \cdot \theta + q \cdot (\theta - 1)] \, \mathrm{d}\left(\mathrm{marg}_\Theta \varphi_1(\tau_1^n)\right)$ and $u_1(\eta^{n,q}, \beta) = 0$, for any $q \in [0, 1]$.

3.3.2 Robustness to higher-order uncertainty about rationality: Rubinstein's Email game

We check now that Osborne and Rubinstein's (1994) version of the Email game is *very* robust to higher-order uncertainty about rationality. We have two players (row and column, named 1 and 2, respectively), two actions available to each of them and two payoff states (α and β) that induce the following payoff structure:

	c	d
a	M, M	1, $-L$
b	$-L$, 1	0, 0

when $\theta = \alpha$

	c	d
a	0, 0	1, $-L$
b	$-L$, 1	M, M

when $\theta = \beta$

where $L > M > 1$. Player 1 knows the true payoff state, and in case it is β, an automatic email is sent to player 2. In case player 2 receives the email, another message is automatically sent to player 1. This process is iterated as long as a player receives an email. Numbers $q, \varepsilon \in (0,1)$ represent the probability of payoff state α and the probability of an email-delivery failure, respectively, and we associate each player's standard type with the number of emails she has sent. Thus, we have sets of standard types $T_1 = T_2 = \mathbb{N} \cup \{0\}$, and belief maps (after some development):

$$\pi_1(t_1)\left[(t_2,\theta)\right] = \begin{cases} (1-q) & \text{if } (t_1,t_2,\theta) = (0,0,\alpha), \\ \frac{1}{2-\varepsilon} & \text{if } (t_1,t_2,\theta) = (t,t-1,\beta) \text{ for some } t > 0, \\ \frac{1-\varepsilon}{2-\varepsilon} & \text{if } (t_1,t_2,\theta) = (t,t,\beta) \text{ for some } t > 0, \\ 0 & \text{otherwise.} \end{cases}$$

$$\pi_2(t_2)\left[(t_1,\theta)\right] = \begin{cases} \frac{(1-q)}{(1-q)+q\varepsilon} & \text{if } (t_1,t_2,\theta) = (0,0,\alpha), \\ \frac{q\varepsilon}{(1-q)+q\varepsilon} & (t_1,t_2,\theta) = (1,0,\beta), \\ \frac{1}{2-\varepsilon} & \text{if } (t_1,t_2,\theta) = (t,t,\beta) \text{ for some } t > 0, \\ \frac{1-\varepsilon}{2-\varepsilon} & \text{if } (t_1,t_2,\theta) = (t,t-1,\beta) \text{ for some } t > 0, \\ 0 & \text{otherwise.} \end{cases}$$

We claim that a and c are the unique ICR^p actions for any type of each player, for any probability $p > \frac{M-1}{L+M-1}$ (note that this lower bound is indeed smaller that $\frac{1}{2}$ for any values of the parameters). In order to see this, notice first that $\mathrm{ICR}_1^{p,\mathscr{T}}(0) = \{a\}$, for any $p > 0$. To see that $\mathrm{ICR}_2^{p,\mathscr{T}}(0) = \{c\}$, for any $p > \frac{M-1}{L+M-1}$, take $\mu \in C_{2,n}^{p,\mathscr{T}}(0)$, where $n \geq 1$ and notice that,

$$\begin{aligned} U_2(\mu,c) - U_2(\mu,d) &\geq \left(\mathrm{marg}_{A_1}\mu\right)[a]L - \left(\mathrm{marg}_{A_1}\mu\right)[b]M + \left(\mathrm{marg}_{A_1}\mu\right)[b] \\ &\geq pL - (1-p)M + (1-p). \end{aligned}$$

Since the last expression is greater than 0, we conclude that $\mathrm{ICR}_{2,n+1}^{p,\mathscr{T}}(0) = \{c\}$, and therefore, that $\mathrm{ICR}_2^{p,\mathscr{T}}(0) = \{c\}$. We proceed now in a inductive way. Let's check first that if $\mathrm{ICR}_2^{p,\mathscr{T}}(t_2) = \{c\}$, then $\mathrm{ICR}_1^{p,\mathscr{T}}(t_2+1) = \{a\}$ for any $t_2 \geq 0$ and any $p > \frac{M-1}{L+M-1}$; to see it, take $\mu \in C_{1,n}^{p,\mathscr{T}}(t_2+1)$, where $n \geq 1$, and note that it must hold that $\left(\mathrm{marg}_{A_2}\mu\right)[c] \geq p$. Then, we have that $U_1(\mu,a) - U_1(\mu,b) \geq pL - (1-p)(M-1)$. Again, since the last expression is bigger than 0, we conclude that $\mathrm{ICR}_{1,n+1}^{p,\mathscr{T}}(t_2+1) = \{a\}$, and therefore, that $\mathrm{ICR}_1^{p,\mathscr{T}}(t_2+1) = \{a\}$. Finally, we check that, if $\mathrm{ICR}_1^{p,\mathscr{T}}(t_1) = \{a\}$, then $\mathrm{ICR}_2^{p,\mathscr{T}}(t_1) = \{c\}$, for any $t_1 \geq 0$ and any $p > \frac{M-1}{L+M-1}$; take $\mu \in C_{2,n}^{p,\mathscr{T}}(t_1)$, where $n \geq 1$ and notice again that it must hold that $\left(\mathrm{marg}_{A_1}\mu\right)[a] \geq p$. Then, we conclude that $\mathrm{ICR}_{2,n+1}^{p,\mathscr{T}}(t_1) = \{c\}$, and thus, due to $U_2(\mu,c) - U_2(\mu,d) \geq pL - (1-p)(M-1)$, that $\mathrm{ICR}_2^{p,\mathscr{T}}(t_1) = \{c\}$. This completes the proof of the claim. We conclude that the Email game exhibits *strong* robustness to higher-order uncertainty about rationality: if rationality is commonly believed with probability $p > \frac{1}{2}$, then (a,c) is the only ICR^p action profile.

3.3.3 Robustness to higher-order uncertainty about rationality: a less robust game

Finally, we elaborate on the example in Section 1.1 to show that robustness to higher-order uncertainty in rationality is not typically as strong as in the previous example. Consider a compact version of Alexei and Polina's (A and P, respectively) game discussed above:

	L	R
T	2, 2	$\frac{2-\theta}{1-\theta}$, 0
B	1, $\frac{-1}{1-\theta}$	$\frac{1}{1-\theta}$, 1

this time, with $\Theta = [0, 1-\varepsilon]$ for some $\varepsilon \in (0,1)$. Standard type structure is given by types $T_A = T_P = \Theta$ and belief maps $\pi_A(t_A)[(t_P,\theta)] = 1_{\{t_A\}}(t_P) \cdot \left(\frac{1}{1-\varepsilon-t_A}\right)$, and $\pi_P(t_P)[(t_A,\theta)] = \pi_A(t_P)[(t_A,\theta)]$ for any $(t_A,t_P,\theta) \in T_A \times T_P \times \Theta$. Obviously, T is the only p-rationalizable action for any type t_A of Alexei and any probability p. We claim that $\mathrm{ICR}_P^p(t_P) = \{L\}$ for any $p > \frac{1+\varepsilon}{1+3\varepsilon}$, and that $\mathrm{ICR}_P^p(1-\varepsilon) = \{L,R\}$ for any $p < \frac{1+\varepsilon}{1+3\varepsilon}$. In order to see this, just note that, for any $t_P < 1-\varepsilon$ and any $\eta \in C_P^p(t_P)$,

$$U_P(\eta,L) - U_P(\eta,R) = 3\eta[T] - 1 - \int_{t_p}^{1-\varepsilon}\left(\frac{1}{1-\theta}\right)(\eta[T] - \eta[T,\theta])\,\mathrm{d}\theta,$$

which is strictly positive, for any $p > \frac{1+\varepsilon}{1+3\varepsilon}$, due to: (i) $3\eta[T] - 1 > 3p - 1$, (ii) $\left(\frac{1}{\varepsilon}\right)(1-p) > \left(\frac{1}{\varepsilon}\right)(1-\eta[T])$, and (iii) $\left(\frac{1}{\varepsilon}\right)(1-\eta[T]) > \int_{t_p}^{1-\varepsilon}\left(\frac{1}{1-\theta}\right)(\eta[T] - \eta[T,\theta])\,\mathrm{d}\theta$. Thus, for any standard type $t_P < 1-\varepsilon$, $\mathrm{ICR}^p(t_P)$ remains equal to $\mathrm{ICR}(t_P)$ for any $p > \frac{1+\varepsilon}{1+3\varepsilon}$, which is close to 1 for ε close to 0, but not so close otherwise. In addition, note that for any $\eta \in C_P^p(1-\varepsilon)$ we have that,

$$U_P(\eta,L) - U_P(\eta,R) = 2\eta[T] - (1-\eta[T])\left(1 + \frac{1}{\varepsilon}\right).$$

Since the value above is negative for any $p < \frac{1+\varepsilon}{1+3\varepsilon}$ and $\eta \in C_P^p(1-\varepsilon)$ where $\eta[T] = p$, we conclude that $R \in \mathrm{ICR}_P^p(1-\varepsilon)$. Thus, if ε is close to 0, we have that $\mathrm{ICR}_P^p(1-\varepsilon) \neq \mathrm{ICR}_P(1-\varepsilon)$ for some p quite close to 1. Note that this formalism, applied to the case in which $\Theta = [0,1)$, that is, non-compact, is enough to easily check that $\mathrm{ICR}_P^p(t_P) = A_p$ for any $p < 1$ and any type t_P. Thus, robustness fails in the non-compact case.

4 Concluding remarks

A. Summary. We introduced interim correlated p-rationalizability (ICR^p) as a solution concept for games with incomplete information and apply it to problems of robustness to higher-order uncertainty about rationality. We prove first that ICR^p captures changes in rational behavior due to departures from common belief in rationality (Theorem 1), which allows us to prove that:

(*i*) Weinstein and Yildiz's (2007) seminal discontinuity result is non-robust (Theorem 2): common belief in rationality is a necessary condition for it to hold. This highlights the fact that higher-order belief restrictions lose their edge under arbitrarily small departures from common belief in rationality.

(*ii*) ICR (Dekel *et al.*, 2007) is robust to higher-order uncertainty. This is an immediate consequence of Theorem 1 and the fact that ICR^p is upper-hemicontinuous on p, and indeed continuous when p equals 1 (Proposition 3), in which case it exactly coincides with ICR.

In addition to this, we also prove some further robustness properties of ICR^p; namely, that it is type-representation independent (Proposition 1) and that it is robust to arbitrarily small misspecification of belief hierarchies (Proposition 2).

B. Alternative topologies on canonical types. The role of the specific topological assumptions we make is obviously, crucial to our results. Both Dekel *et al.* (2006) and Chen *et al.* (2010) present detailed discussions on topological concerns that are beyond the scope of this paper. However, the topologies considered in those papers are all finer than the one we make use of, and thus,[31] for any of such topologies:

(*i*) The non-robustness of Weinstein and Yildiz's discontinuity to higher-order uncertainty about rationality presented in Theorem 2 holds: just note that for any pair of topologies $\mathcal{O}_1$ and $\mathcal{O}_2$ where $\mathcal{O}_2$ is finer than $\mathcal{O}_1$, any sequence which is convergent w.r.t. $\mathcal{O}_2$ will *a fortiori* be convergent w.r.t. $\mathcal{O}_1$, and thus, the hypothesis in the theorem apply.

(*ii*) The robustness of ICR^p to higher-order uncertainty about rationality derived from Theorem 1 and Proposition 3 holds; despite being somewhat tedious, it is routine to check that topological assumptions in the proofs of the results, are still satisfied.

[31] Namely, these authors introduce the uniform-strategic, the strategic and the uniform-weak topologies.

C. PRIVATE INFORMATION. Some authors identify private information with knowledge of some feature of the payoff state (see Battigalli *et al.*, 2011 or Penta, 2012, or Bergemann and Morris, 2005, for an example in mechanism design literature). This is done writing $\Theta = \Theta_0 \times \prod_{i\in I} \Theta_i$, and structurally assuming that at each payoff state θ, each player i is right about the i^{th} coordinate θ_i. In case we were interested in having this interpretation of private information made explicit, we need first to take event $PI \cap CB(PI)$, where $PI = \bigcap_{i\in I} PI_i$ and for each player i, $PI_i = \left\{\omega \in \Omega \left| e_{i,1}(\omega) = 1_{\{\theta_i(\omega)\}}\right.\right\}$. That is, consider the event in which player i is informed about the i-th coordinate of the payoff state. Define:

$$\mathcal{T}^0_{i,PI} = \bigcup_{\theta_i\in\Theta_i} \left\{\tau_i \in \mathcal{T}_i \left| \tau_{i,1} = 1_{\{\theta_i\}}\right.\right\},$$
$$\mathcal{T}^{n+1}_{i,PI} = \left\{\tau_i \in \mathcal{T}^n_{i,PI} \left| \varphi_i(\tau_i)\left[\Theta \times \mathcal{T}^n_{-i,PI}\right] = 1\right.\right\},$$

for any $n \in \mathbb{N}$, and set $\mathcal{T}_{i,PI} = \bigcap_{n\geq 0} \mathcal{T}^n_{i,PI}$. Each standard hierarchy in $\mathcal{T}_{i,PI}$ represents common belief in each player being right about her corresponding coordinate of the payoff state, and it is routine to check that, for any $\omega \in PI_i \cap CB(PI)$ and any player i, it holds that $q_i(e_i(\omega)) \in \mathcal{T}_{i,PI}$. Thus, we obtain the following epistemic characterization in this case:

$$\bigcup_{\tau_i\in\mathcal{T}_{i,PI}} \mathrm{ICR}^p_i(\tau_i) = \mathrm{Proj}_{A_i}\left(R_i \cap CB^p(R) \cap PI_i \cap CB(PI)\right).$$

for any player i and probability p. While making use of a product structure can turn out to be insightful and can provide clearer intuition in some circumstances, we opted for the non-restricted structure for the sake of notational simplicity and the fact that keeping track of the additional information does not play a substantial role in our specific analysis.

A PROOFS

Note first that it is easy to provide a fixed-point characterization of the solution concepts introduced so far in a trivial way. We explicitly check this for ICR^p, but a similar argument applies both to ICR and ICSR. Denote:

$$\begin{aligned}\mathrm{ICR}^{p,*}_i(\tau_i) &= \{a_i \in A_i \,| a_i \in BR_i(\eta) \text{ for some } \eta \in C^p_i(\tau_i)\} \text{ and,}\\ C^p_i(\tau_i) &= \left\{\eta \in C^p_{i,0}(\tau_i) \left| \eta\left[\Theta \times \mathrm{Graph}\left(\mathrm{ICR}^p_{-i}\right)\right] \geq p\right.\right\}.\end{aligned}$$

We claim that $\mathrm{ICR}^{p,*}_i(\tau_i) = \mathrm{ICR}^p_i(\tau_i)$. For the right inclusion, note that since, for any $n \in \mathbb{N}$, we have that $\mathrm{ICR}^p_{-i}(\tau_{-i}) \subseteq \mathrm{ICR}^p_{-i,n}(\tau_{-i})$, for any standard hierarchy τ_{-i}, then, it holds that $C^p_i(\tau_i) \subseteq C^p_{i,n}(\tau_i)$ too, and thus, that $\mathrm{ICR}^{p,*}_i(\tau_i) \subseteq \mathrm{ICR}^p_i(\tau_i)$. For the left inclusion, take sequence $(\eta_n)_{n\in\mathbb{N}} \in \prod_{n\in\mathbb{N}} C^p_{i,n}(\tau_i)$ such that $a_i \in BR_i(\eta_n)$ for any $n \in \mathbb{N}$; the fact that $\Delta\left(\mathcal{T}_{-i} \times A_{-i} \times \Theta\right)$ is compact implies that there is some convergent subsequence $(\eta_{n_m})_{m\in\mathbb{N}} \subseteq (\eta_n)_{n\in\mathbb{N}}$ with limit $\eta \in \bigcap_{n\geq 0} C^p_{i,n}(\tau_i)$, which, therefore, satisfies $\eta\left[\Theta \times \text{Graph}\left(\text{ICR}^p_{-i}\right)\right] \geq p$.

A.1 Elementary properties of ICR^p

A.1.1 Epistemic foundation of ICR^p

Theorem 1 (Epistemic foundation of ICR^p). *Let $\mathscr{G}$ be a game with incomplete information, and p, a positive probability. Then, interim correlated p-rationalizability characterizes rationality and common p-belief in rationality; i.e., for any player i and any standard hierarchy τ_i,*

$$\text{ICR}^p_i(\tau_i) = \text{Proj}_{A_i}\left(R_i \cap B_i(CB^p(R)) \cap [q_i = \tau_i]\right).$$

Proof. We prove the first statement and then establish how the second statement is proved in an analogous way. Fix positive probability p. For each player i, let $X_{i,0} = R_i$, and, for each $n \in \mathbb{N}$, let $X_{i,n} = X_{i,n-1} \cap B^p_i(X_{-i,n-1})$, where $X_{-i,n-1} = \bigcap_{j\neq i} X_{j,n-1}$. It is immediate that, for any $n \in \mathbb{N}$, it holds that (i) $X_{i,n} = R_i \cap B^p_i(B^{n-1,p}(R))$, (ii) $X_{i,n-1} \subseteq X_{i,n}$ and (iii) $X_{i,n}$ is closed. Our objective is to prove that, for any $n \geq 0$,

$$\text{Proj}_{\Omega_i} X_{i,n} = \text{Graph}\left(\text{ICR}^p_{i,n+1} \circ q_i^{-1}\right).$$

We begin with the right inclusion. Let's proceed by induction: take pair $(e_i, a_i) \in \text{Proj}_{\Omega_i} X_{i,0}$ and define conjecture $\eta \in \Delta(\mathcal{T}_{-i} \times A_{-i} \times \Theta)$ as follows: $(\tau_{-i}, a_{i-}, \theta) \mapsto \psi_i(e_i)\left[q^{-1}_{-i}(\tau_{-i}) \times \{(a_{-i}, \theta)\}\right]$. It is immediate that $\text{marg}_{A_{-i}\times\Theta}\eta = e_{i,1}$, and also that $\text{marg}_{\mathcal{T}_{-i}\times\Theta}\eta = \text{marg}_{\mathcal{T}_{-i}\times\Theta}\varphi_i(q_i(e_i))$. Thus, we conclude that $a_i \in \text{ICR}^p_{i,1}(q_i(e_i))$. Now, let $n \in \mathbb{N}$ be such that $\text{Proj}_{\Omega_i} X_{i,k} \subseteq \text{Graph}\left(\text{ICR}^p_{i,k+1} \circ q_i^{-1}\right)$, for any $k < n$ and any $i \in I$. Let $(e_i, a_i) \in \text{Proj}_{\Omega_i} X_{i,n}$ and conjecture $\eta \in \Delta(\mathcal{T}_{-i} \times A_{-i} \times \Theta)$, where $(\tau_{-i}, a_{i-}, \theta) \mapsto \psi_i(e_i)\left[q^{-1}_{-i}(\tau_{-i}) \times \{(a_{-i}, \theta)\}\right]$. Again, both $\text{marg}_{A_{-i}\times\Theta}\eta = e_{i,1}$, and $\text{marg}_{\mathcal{T}_{-i}\times\Theta}\eta = \text{marg}_{\mathcal{T}_{-i}\times\Theta}\varphi_i(q_i(e_i))$ are immediate. In addition,

$$\begin{aligned}\eta\left[\Theta \times \text{Graph}\left(\text{ICR}^p_{-i,n}\right)\right] &= \psi_i(e_i)\left[\Theta \times \text{Graph}\left(\text{ICR}^p_{-i,n} \circ q^{-1}_{-i}\right)\right] \\ &\geq \psi_i(e_i)\left[\Theta \times \text{Proj}_{\Omega_{-i}} X_{-i,n-1}\right].\end{aligned}$$

Since $(e_i, a_i) \in \text{Proj}_{\mathcal{E}_i\times A_i} B^p_i(X_{-i,n})$, we know that $\psi_i(e_i)\left[\Theta \times \text{Proj}_{\Omega_{-i}} X_{-i,n}\right] \geq p$. Thus, $a_i \in \text{ICR}^{p,\mathscr{T}}_{i,n+1}(q_i(e_i))$ and we conclude the right inclusion of the statement.

Before going on with the proof of the left inclusion, for any $i \in I$, let correspondence $\Phi_{i,0} : \mathcal{T}_i \times A_i \rightrightarrows \Omega_i$ be given by $(\tau_i, a_i) \mapsto q_i^{-1}(\tau_i) \times \{a_i\}$, and for any $n \in \mathbb{N}$, let correspondence $\Phi_{i,n} : \text{Graph}\left(\text{ICR}^p_{i,n}\right) \rightrightarrows \Omega_i$ be given by,

$$(\tau_i, a_i) \mapsto \left\{e_i \in q_i^{-1}(\tau_i) \,\middle|\, (e_i, a_i) \in \text{Proj}_{\Omega_i} X_{i,n-1}\right\} \times \{a_i\}.$$

Note that $\Phi_{i,n+1}\left(\text{Graph}\left(\text{ICR}^p_{i,n+1}\right)\right) \subseteq \text{Proj}_{\Omega_i} X_{i,n}$. In addition, each correspondence $\Phi_{i,n}$ has closed graph. In order to see this, consider sequence $(\tau_i^m, a_i^m; e_i^m, a_i^m)_{m\in\mathbb{N}} \subset$

Graph$(\Phi_{i,n})$ with limit $(\tau_i, a_i; e_i, a_i)$. We need to check that $(e_i, a_i) \in \Phi_{i,n}(\tau_i, a_i)$. Since sets $[q_i = \tau_i]$ and $X_{i,n}$ are closed, it follows that their intersection is closed, and, since Ω_i is compact and Hausdorff, map Proj_{Ω_i} is closed. This implies closeness of set $\left\{e_i \in \mathcal{E}_i \,\middle|\, (e_i, a_i) \in \mathrm{Proj}_{\Omega_i}([q_i = \tau_i] \cap X_{i,n-1})\right\}$. Now take sequence $(\tau_i^m, a_i^m; e_i^m, a_i^m)_{m\in\mathbb{N}}$; then there is some sequence $(\tau_i^m, a_i^m; \omega^m)_{m\in\mathbb{N}}$ with $e_i(\omega^m) = e_i^m$ and $a_i(\omega^m) = a_i^m$ such that $\omega^m \in [q_i = \tau_i] \cap X_{i,n-1}$, for every m, hence also $q_i(e_i(\omega^m)) = \tau_i^m$. Because $X_{i,n-1}$ is closed and fixed (over the domain of $\Phi_{i,n}$), clearly the limit ω of ω^m is in $X_{i,n-1}$. To see that it is also in $[q_i = \tau_i]$, notice again that q_i and e_i are continuous, so that $q_i(e_i(\omega)) = \tau_i$, since by construction and by continuity $\lim_{m\to\infty} \tau_i^m = \lim_{m\to\infty} q_i(e_i(\omega^m)) = q_i(e_i(\omega))$. Hence $\omega \in [q_i = \tau_i] \cap X_{i,n-1}$ and so, by continuity, we have $\mathrm{Proj}_{\mathcal{E}_i \times A_i}(\omega) = (e_i, a_i)$ and indeed $(e_i, a_i) \in \Phi_{i,n}(\tau_i, a_i)$.

For the proof of the left inclusion, we proceed again by induction. $\Phi_{i,0}$ is a non-empty correspondence with closed graph. Thus, it is weakly measurable, and therefore, we know because of the Kuratowski-Ryll Nardzewski Selection Theorem, that there exists some measurable map $\phi_{i,0} : \mathcal{T}_i \times A_i \to \Omega_i$ such that $\phi_{i,0}(\tau_i, a_i) \in \Phi_{i,0}(\tau_i, a_i)$, for any $(\tau_i, a_i) \in \mathcal{T}_i \times A_i$. We denote $\phi_{-i,0} = (\phi_{j,0})_{j\neq i}$. Now, take $(\tau_i, a_i) \in \mathrm{Graph}\left(\mathrm{ICR}^p_{i,1}\right)$, and conjecture $\eta \in C^p_{i,0}(\tau_i)$ that rationalizes a_i. We define $\mu(\eta) \in \Delta(\Omega_{-i} \times \Theta)$ as follows: $(e_{-i}, a_{-i}, \theta) \mapsto \eta\left[\phi^{-1}_{-i,0}(e_{-i}, a_{-i}) \times \{\theta\}\right]$. Then, for $e_i = \psi_i^{-1}(\mu(\eta))$ it holds that $e_{i,1} = \mathrm{marg}_{A_{-i}\times\Theta}\eta$, and therefore, that $a_i \in BR_i(e_{i,1})$, *i.e.*, that $(e_i, a_i) \in \mathrm{Proj}_{\Omega_i} X_{i,0}$. Note that, in particular, $(e_i, a_i) \in \Phi_{i,1}(\tau_i, a_i)$, so that $\Phi_{i,1}$ is a non-empty correspondence.

Let $n \in \mathbb{N}$ such that $\mathrm{Graph}\left(\mathrm{ICR}^{p,\mathscr{T}}_{i,k+1}\right) \subseteq \mathrm{Proj}_{\Omega_i} X_{i,k}$ and $\Phi_{i,k+1}$ is non-empty for any $k < n$, $(\tau_i, a_i) \in \mathrm{Graph}\left(\mathrm{ICR}^p_{i,n+1}\right)$, and conjecture $\eta \in C^p_{i,n}(\tau_i)$ that rationalizes a_i. We define $\mu(\eta) \in \Delta(\Omega_{-i} \times \Theta)$ as follows: $(e_{-i}, a_{-i}, \theta) \mapsto \eta\left[\phi^{-1}_{-i,n}(e_{-i}, a_{-i}) \times \{\theta\}\right]$ where each $\phi_{j,n}$ is again a measurable selector whose existence is assured by the Kuratowski-Ryll Nardzewski Selection Theorem. Then, for epistemic type $e_i = \psi_i^{-1}(\mu(\eta))$ it holds that $a_i \in BR_i(e_{i,1})$, because $e_{i,1} = \mathrm{marg}_{A_{-i}\times\Theta}\eta$. The fact that,

$$\psi_i(e_i)\left[\Theta \times \mathrm{Proj}_{\Omega_i} X_{-i,n}\right] \geq \eta\left[\Theta \times \mathrm{Graph}\left(\mathrm{ICR}^p_{-i,n-1}\right)\right] \geq p,$$

lets us conclude that $(e_i, a_i) \in \mathrm{Proj}_{\Omega_i} X_{i,n}$ and also that $(e_i, a_i) \in \Phi_{i,n+1}(\tau_i, a_i)$, so that $\Phi_{i,n+1}$ is non-empty.

Thus, we conclude that $\mathrm{Proj}_{\Omega_i} X_{i,n} = \mathrm{Graph}\left(\mathrm{ICR}^p_{i,n+1} \circ q_i^{-1}\right)$ for any $n \geq 0$. The fact that $\prod_{i\in I} \mathrm{Proj}_{A_i} \bigcap_{n\geq 0} X_{i,n} = \mathrm{Proj}_A(R \cap CB^p(R))$ finishes the proof of the first statement of the theorem. The proof of the second statement is completed in an analogous way, introducing the following change: define now, $X_{i,0} = \mathrm{ICR}_i \cap [q_i = \tau_i]$. ■

A.1.2 Robustness to type-representation

Proposition 1 (Type-representation independence). *Let Γ be a Bayesian game. Then, for any player i, any standard type t_i, any positive probability p and any non-negative integer n, it holds that* $\mathrm{ICR}^{p,\mathscr{T}}_{i,n}(t_i) = \mathrm{ICR}^p_{i,n}(\tau_i(t_i))$.

Proof. Recall first that we know from (Battigalli *et al.*, 2011, Remark 1), that for any player i and any standard type t_i it holds that, for any standard type t_i and any event $E \in \mathscr{B}(\mathcal{T}_{-i} \times \Theta)$,

$$\varphi_i(\tau_i(t_i))[E] = \pi_i(t_i)[\{(t_{-i},\theta) \in T_{-i} \times \Theta \,|(\tau_{-i,}(t_{-i}),\theta) \in E\}]. \tag{2}$$

Now, the case $n = 0$ is trivial. Let's proceed by induction and assume that $n \in \mathbb{N}$ is such that, for any $k < n$, we have that, $\mathrm{ICR}^{p,\mathscr{T}}_{j,k}(t_j) = \mathrm{ICR}^p_{j,k}(\tau_j(t_j))$ for any $j \in I$ and any $t_j \in T_j$. Our aim is to prove the equality for n. Fix $i \in I$ and $t_i \in T_i$. We begin with the right inclusion. Take $a_i \in \mathrm{ICR}^{p,\mathscr{T}}_{i,n}(t_i)$ and $\mu \in C^{p,\mathscr{T}}_{i,n-1}(t_i)$, a belief that rationalizes a_i. We define $\eta(\mu) \in \Delta(\mathcal{T}_{-i} \times A_{-i} \times \Theta)$ as: let $\eta(\mu)[(\tau_{-i}, a_{-i}, \theta)] = \mu\left[\tau^{-1}_{-i}(\tau_{-i}) \times \{(a_{-i},\theta)\}\right]$ for any $(\tau_{-i}, a_{-i}, \theta) \in \mathcal{T}_{-i} \times A_{-i} \times \Theta$.[32] Then: (i) it holds that $a_i \in BR_i(\eta(\mu))$, because $\mathrm{marg}_{A_{-i}\times\Theta}\eta(\mu) = \mathrm{marg}_{A_{-i}\times\Theta}\mu$, and (ii) we have that, $\mathrm{marg}_{\mathcal{T}_{-i}\times\Theta}\eta(\mu)[(\tau_{-i},\theta)] = \pi_i(t_i)\left[\tau^{-1}_{-i}(\tau_{-i}) \times \{\theta\}\right]$ for any $(\tau_{-i},\theta) \in \mathcal{T}_{-i} \times \Theta$, and thus, due to equation (2), it holds that $\mathrm{marg}_{\mathcal{T}_{-i}\times\Theta}\eta(\mu) = \varphi_i(\tau_i(t_i))$. To finish the proof of this inclusion, note that we have,

$$\begin{aligned}
&\eta(\mu)\left[\Theta \times \mathrm{Graph}\left(\mathrm{ICR}^p_{-i,n-1}\right)\right] \geq \\
&\quad \geq \mu\left[\Theta \times \left\{(t_{-i},a_{-i}) \in T_{-i} \times A_{-i}\,\middle|\, a_{-i} \in \mathrm{ICR}^p_{-i,n-1}(\tau_{-i}(t_{-i}))\right\}\right] \\
&\quad = \mu\left[\Theta \times \mathrm{Graph}\left(\mathrm{ICR}^{p,\mathscr{T}}_{-i,n-1}\right)\right] \geq p.
\end{aligned}$$

Let's continue with the left inclusion. Let $a_i \in \mathrm{ICR}^p_{i,n}(\tau_i(t_i))$ and $\eta \in C^p_{i,n-1}(\tau_i(t_i))$, a belief that rationalizes a_i. We now define belief $\mu(\eta) \in \Delta(T_{-i} \times A_{-i} \times \Theta)$ as follows: let $\mu(\eta)[(t_{-i}, a_{-i}, \theta)] = \delta^\theta_i(t_{-i})\,\eta[(\tau_{-i}(t_{-i}), a_{-i}, \theta)]$ for any $(t_{-i}, a_{-i}, \theta) \in T_{-i} \times A_{-i} \times \Theta$, where we have denoted $\delta^\theta_i(t_{-i}) = \pi_i(t_i)\left[(t_{-i},\theta)\,\middle|\,\tau^{-1}_{-i}(\tau_i(t_{-i})) \times \{\theta\}\right]$. Note that for any $\theta \in \Theta$ and any $t_{-i} \in T_{-i}$ we have $\delta^\theta_i\left(\tau^{-1}_{-i}(\tau_{-i}(t_{-i}))\right) = 1$, and thus, $\mu(\eta)\left[\tau^{-1}_{-i}(\tau_{-i}(t_{-i})) \times \{(a_{-i},\theta)\}\right] = \eta[(\tau_i(t_{-i}), a_{-i}, \theta)]$, for any $(t_{-i}, a_{-i}, \theta) \in T_{-i} \times A_{-i} \times \Theta$. From this equality, first, we have that $\mathrm{marg}_{A_{-i}\times\Theta}\mu(\eta) = \mathrm{marg}_{A_{-i}\times\Theta}\eta$, and thus, $a_i \in BR_i(\mu(\eta))$; and second, $\mu(\eta)\left[\Theta \times \tau^{-1}_{-i}(\tau_{-i}(t_{-i})) \times \{a_{-i}\}\right] = \eta[\Theta \times \{(\tau_i(t_{-i}), a_{-i})\}]$, and thus, it follows from the induction hypothesis that $\mu(\eta)\left[\Theta \times \mathrm{Graph}\left(\mathrm{ICR}^{p,\mathscr{T}}_{-i,n-1}\right)\right] = \eta\left[\Theta \times \mathrm{Graph}\left(\mathrm{ICR}^p_{-i,n-1}\right)\right] \geq p$.[33] To see that $\mathrm{marg}_{T_{-i}\times\Theta}\mu(\eta) = \pi_i(t_i)$, check that,

$$\begin{aligned}
\mathrm{marg}_{T_{-i}\times\Theta}\mu(\eta)[(t_{-i},\theta)] &= \delta^\theta_{-i}(t_{-i})\,\eta[A_{-i} \times \{(\tau_{-i}(t_{-i}),\theta)\}] \\
&= \delta^\theta_{-i}(t_{-i})\,\varphi_i(\tau_i(t_i))[(\tau_{-i}(t_{-i}),\theta)] \\
&= \delta^\theta_{-i}(t_{-i})\,\pi_i(t_{-i})\left[\tau^{-1}_{-i}(\tau_{-i}(t_{-i})) \times \{\theta\}\right] \\
&= \pi_i(t_i)[(t_{-i},\theta)].
\end{aligned}$$

for any $(t_{-i},\theta) \in T_{-i} \times \Theta$. ■

[32] The continuity of each τ_j assures that the definition is correct; it is straightforward how to extend the pointwise definition to the σ-algebra.

[33] For this step note also that if $\tau_{-i} \notin \tau_{-i}(T_{-i})$, then $\mathrm{marg}_{\mathcal{T}_{-i}}\eta[\tau_{-i}] = 0$.

A.1.3 Robustness to higher-order uncertainty about payoffs

LEMMA 1. *Let $\mathscr{G}$ be a game with incomplete information, i, a player, and E_i, a closed subset of $\mathcal{E}_i$ such that $q_i(E_i) = \mathcal{T}_i$. Then, correspondence $\Sigma_i^{E_i}$ given by,*

$$\begin{array}{rccc} \Sigma_i^{E_i}: & \mathcal{T}_i & \rightrightarrows & A_i \\ & \tau_i & \mapsto & \operatorname{Proj}_{A_i}\left(R_i \cap \left(\mathcal{E}_{-i} \times \left(q_i^{-1}(\tau_i) \cap E_i\right) \times A \times \Theta\right)\right), \end{array}$$

is upper-hemicontinuous.

Proof. It follows from Proposition 1 by Piermont and Zuazo-Garin (2015). Following their notation and terminology, just model $\mathscr{G}$ as sequential game $\mathscr{S} = \langle \mathscr{T}, \mathscr{P} \rangle$, where we have sequential structure $\mathscr{T} = \left\langle I, (A_i)_{i\in I}, H, Z \right\rangle$ with $H = \{\emptyset\}$ and $Z = \{(\emptyset, a) \mid a \in A\}$, and payoff structure $\mathscr{P} = \left\langle \Theta, (U_i)_{i\in I} \right\rangle$ with $U_i(z, \theta) = u_i(a, \theta)$ for any $i \in I$ and any $z = (\emptyset, a)$ where $a \in A$. Then, just apply the epistemic characterization in Theorem 1 and note that $B_i^p(CB^p(R))$ is closed. ■

Lemma 1 provides a useful insight into the relation between upper-hemicontinuity and topological properties of the epistemic restrictions.[34] It should be noted that it does not crucially depend on the fact that we endow spaces of probability measures with the weak$*$ topology, but rather, on $\mathcal{E}_i$ being compact, and A_i, Hausdorff.[35] Therefore, the result holds for any topology on spaces of probability measures that implies $\mathcal{E}_i$ being compact. Condition $q_i(E_i) = \mathcal{T}_i$ should be read as the fact that E_i is a set of restrictions which only involves uncertainty concerning choices, not uncertainty concerning payoff states. The lemma straightforwardly implies the upper-hemicontinuity of ICR^p on standard hierarchies, which can be understood as a robustness property of ICR^p.

PROPOSITION 2 (Robustness to uncertainty about payoffs). *Let $\mathscr{G}$ be a game with incomplete information, and p, a positive probability. Then, for any player i, correspondence $\mathrm{ICR}_i^p : \mathcal{T}_i \rightrightarrows A_i$ is upper-hemicontinuous.*

Proof. Note first, that $B_i^p\left(CB^p(R)\right) = \Omega_{-i} \times A_i \times \Theta \times \operatorname{Proj}_{\mathcal{E}_i} CB^p(R)$ is closed, and second, that $\mathrm{ICR}_i^p = \Sigma_i^{E_i}$ for $E_i = B_i^p\left(CB^p(R)\right)$. Then, it follows directly from Theorem 1 and Lemma 1. ■

A.1.4 Robustness to higher-order uncertainty about rationality

PROPOSITION 3 (Continuity properties on p). *Let $\mathscr{G}$ be a game with incomplete information. Then, for any player i and any standard hierarchy τ_i, correspondence given by $p \mapsto \mathrm{ICR}_i^p(\tau_i)$ is upper-hemicontinuous, continuous at $p = 1$, and satisfies that $1 \mapsto \mathrm{ICR}_i(\tau_i)$.*

[34] Not to be confused with topological properties of the graph of the correspondence. The relation between the latter and continuity properties of the correspondence is well known.

[35] This last requirement is satisfied *a fortiori*, if utility maps are continuous and any pair of actions of any player delivers different payoffs at some payoff state, or against some partial profile of the opponents'.

Proof. We proceed in four steps:

CLAIM 1. *For any $i \in I$, any $\tau_i \in \mathcal{T}_i$, any $n \geq 0$ and any $p \in (0,1]$, $C^p_{i,n}(\tau_i)$ is closed.* Take convergent sequence $(\eta_k)_{k\in\mathbb{N}} \subseteq C^p_{i,n}(\tau_i)$ with limit η. It is obvious that $\text{marg}_{\mathcal{T}_{-i}\times\Theta}\eta = \varphi_i(\tau_i)$. Since $\eta_k\left[\Theta \times \text{Graph}\left(\text{ICR}^p_{-i,n-1}\right)\right] \geq p$ for any $k \in \mathbb{N}$, we have that $\eta\left[\Theta \times \text{Graph}\left(\text{ICR}^p_{-i,n-1}\right)\right] = \lim_{k\to\infty}\eta_k\left[\Theta \times \text{Graph}\left(\text{ICR}^p_{-i,n-1}\right)\right] \geq p$, and thus, that $\eta \in C^p_{i,n}(\tau_i)$.. ★

CLAIM 2. *For any $i \in I$, any $\tau_i \in \mathcal{T}_i$ and any $n \geq 0$, correspondence $\Gamma^{\tau_i}_n$ given by $p \mapsto C^p_{i,n}(\tau_i)$ is continuous at any $p \in (0,1]$.* For upper-hemicontinuity, take convergent sequence $(p_k,\eta_k)_{k\in\mathbb{N}}$ with limit (p,η), where $\eta_k \in C^{p_k}_{i,n}(\tau_i)$ for any $k \in \mathbb{N}$. Then, it holds that $\eta_k\left[\Theta \times \text{Graph}\left(\text{ICR}^{p_k}_{-i,n-1}\right)\right] \geq p_k$ for any $k \in \mathbb{N}$, and thus, that $\eta\left[\Theta \times \text{Graph}\left(\text{ICR}^p_{-i,n-1}\right)\right] = \lim_{k\to\infty}\eta_k\left[\Theta \times \text{Graph}\left(\text{ICR}^{p_k}_{-i,n-1}\right)\right] \geq \lim_{k\to\infty}p_k = p$. Since it is obvious that $\text{marg}_{\mathcal{T}_{-i}\times\Theta}\eta = \varphi_i(\tau_i)$, we conclude that $\eta \in C^p_{i,n}(\tau_i)$. For lower-hemicontinuity, take convergent $(p_k)_{k\in\mathbb{N}} \subseteq (0,1]$ with limit $p \in (0,1]$, and $\eta \in C^p_{i,n}(\tau_i)$. Then, for each $k \in \mathbb{N}$, and each $(\tau_{-i},a_{-i},\theta) \in \mathcal{T}_{-i} \times A_{-i} \times \Theta$, let:

$$\eta_k\left[(\tau_{-i},a_{-i},\theta)\right] = \begin{cases} \left(\frac{p_k}{p^*}\right)\eta\left[(\tau_{-i},a_{-i},\theta)\right] & \text{if } a_{-i} \in \text{ICR}^p_{-i,n-1}(\tau_{-i}), \\ \delta^{(\tau_{-i},\theta)}\eta\left[(\tau_{-i},a_{-i},\theta)\right] & \text{otherwise,} \end{cases}$$

where we have $\delta^{(\tau_{-i},\theta)} = \left(1 + \frac{1-\left(\frac{p_k}{p^*}\right)\eta\left[\text{ICR}^p_{-i,n-1}(\tau_{-i})\times\{(\tau_{-i},\theta)\}\right]}{\varphi_i(\tau_i)[(\tau_{-i},\theta)]-\eta\left[\text{ICR}^p_{-i,n-1}(\tau_{-i})\times\{(\tau_{-i},\theta)\}\right]}\right)$, and we set probability $p^* = \mu\left[\Theta \times \text{Graph}\left(\text{ICR}^p_{-i,n-1}\right)\right] \geq p$. It is immediate that η_k is a well-defined probability measure. It is easy to check that (i) $\text{marg}_{\mathcal{T}_{-i}\times\Theta}\eta_k = \varphi_i(\tau_i)$, and (ii) $\eta_k\left[\Theta \times \text{Graph}\left(\text{ICR}^{p_k}_{-i,n-1}\right)\right] = \eta_k\left[\Theta \times \text{Graph}\left(\text{ICR}^p_{-i,n-1}\right)\right] = p_k$, and therefore, $\eta_k \in C^{p_k}_{i,n}(\tau_i)$ for every $k \in \mathbb{N}$. The fact that $\lim_{k\to\infty}\eta_k = \eta$ completes the proof. ★

CLAIM 3. *For any $n \geq 0$, any $p \in (0,1]$, any $i \in I$ and any $\tau_i \in \mathcal{T}_i$, $\text{ICR}^p_{i,n}(\tau_i)$ is a closed subset of A_i. Thus, correspondence $p \mapsto \text{ICR}^p_{i,n}(\tau_i)$ is compact-valued.* Let convergent sequence $\left(a^k_i\right)_{k\in\mathbb{N}} \subseteq \text{ICR}^p_{i,n}(\tau_i)$ with limit a_i. We know that there is some sequence $\left(\eta^k\right)_{k\in\mathbb{N}} \subseteq C^p_{i,n-1}(\tau_i)$ such that $a^k_i \in BR_i\left(\eta^k\right)$ for any $k \in \mathbb{N}$. From compactness of $C^p_{i,n-1}(\tau_i)$ (closed, due to Claim 1, and also subset of a compact space) we know that there exists a convergent subsequence $\left(\eta^{k_l}\right)_{l\in\mathbb{N}} \subseteq \left(\eta^k\right)_{k\in\mathbb{N}}$ with limit $\eta \in C^p_{i,n-1}(\tau_i)$. Obviously, subsequence $\left(a^{k_l}_i\right)_{l\in\mathbb{N}}$ converges to a_i, and thus, due to the best response correspondence being compact-valued and upper-hemicontinuous, $a_i \in BR_i(\eta)$. Thus, $a_i \in \text{ICR}^p_{i,n}(\tau_i)$, and therefore, $p \mapsto \text{ICR}^p_{i,n}(\tau_i)$ is compact-valued. ★

PROOF OF THE PROPOSITION. The claims above make the proof of the proposition straightforward. As mentioned before, it is obvious that $\text{ICR}_i(\tau_i) = \text{ICR}^1_i(\tau_i)$ for any $i \in I$ and any $\tau_i \in \mathcal{T}_i$, so we only have to prove upper-hemicontinuity. First, note that $\Gamma^{\tau_i} = \bigcap_{n\in\mathbb{N}}\Gamma^{\tau_i}_n$, where for each $n \in \mathbb{N}$, $\Gamma^{\tau_i}_n : (0,1] \rightrightarrows A_i$ is given by $p \mapsto \text{ICR}^p_{i,n}(\tau_i)$. Note additionally that $\Gamma^{\tau_i}_n$ is the composition of correspondences $p \mapsto C^p_{i,n-1}(\tau_i)$ and

$\mu \mapsto BR_i(\mu)$. We checked above that $p \mapsto C^p_{i,n-1}(\tau_i)$ is upper-hemicontinuous, and we know that so is $\mu \mapsto BR_i(\mu)$. Thus, by Theorem 17.23 in Aliprantis and Border (1999), $\Gamma^{\tau_i}_n$ is upper-hemicontinuous too. We saw in the Claim 2 that $\Gamma^{\tau_i}_n$ is closed, and thus, since A_i is compact, $\Gamma^{\tau_i}_n$ is both closed and compact valued. Hence, since $(0,1]$ is regular, by Theorem 17.25 in Aliprantis and Border (1999), that obtain that Γ^{t_i} is upper-hemicontinuous. To check continuity at $p=1$ it suffices to show that Γ^{τ_i} is lower-hemicontinuous at $p=1$. That is, we need to show (see Aliprantis and Border 1999, Def. 17.2) that for any open subset $U \subseteq A_i$ such that $\Gamma^{\tau_i}(1) \cap U \neq \emptyset$, there exists a neighborhood $V \subseteq (0,1]$ of $p=1$ such that if $p' \in V$, then $\Gamma^{\tau_i}(p') \cap U \neq \emptyset$. This follows immediately from the fact that the set of actions contained in Γ^{τ_i} increases as p decreases. In particular, $\Gamma^{\tau_i}(1) \subseteq \Gamma^{\tau_i}(p')$ for any $p' \in V$, and hence $\Gamma^{\tau_i}(p') \cap U \neq \emptyset$ for any $p' \in V$. Since Γ^{τ_i} is upper-hemicontinuous on all of $(0,1]$, combining this with lower-hemicontinuity at $p=1$ shows that Γ^{τ_i} is indeed continuous at $p=1$. ■

A.2 Non-robustness of Weinstein and Yildiz's discontinuity

Theorem 2 (Non-robustness of Weinstein and Yildiz's discontinuity). *Let $\mathscr{G}$ be a game with incomplete information and finite set of action profiles. Then for any standard hierarchy τ_i, any $p<1$ and any sequence $(\tau_i^n)_{n\in\mathbb{N}}$ converging to τ_i, there is some $n_p \in \mathbb{N}$ such that $\mathrm{ICSR}_i(\tau_i) \subseteq \mathrm{ICR}^p_i(\tau^m_i)$ for any $m \geq n_p$.*

Proof. We proceed in four steps:

Claim 1. *For any convergent sequence of standard hierarchies $(\tau_i^n)_{n\in\mathbb{N}}$ with limit τ_i and any $\eta \in \Delta(\mathcal{T}_{-i}\times A_{-i}\times\Theta)$ such that $\mathrm{marg}_{\mathcal{T}_{-i}\times\Theta}\eta = \varphi_i(\tau_i)$, there exists a convergent sequence of beliefs $(\eta^n)_{n\in\mathbb{N}} \subseteq \Delta(\mathcal{T}_{-i}\times A_{-i}\times\Theta)$ with limit η, and natural $M\in\mathbb{N}$, such that $\mathrm{marg}_{\mathcal{T}_{-i}\times\Theta}\eta^m = \varphi_i(\tau_i^m)$, for any $m\geq M$.* For continuity, we have that since $\tau_i^n \to \tau_i$, then $\varphi_i(\tau_i^n)\to\varphi_i(\tau_i)$, and thus, if we denote $S=\mathrm{supp}\varphi_i(\tau_i)$, there is some $M\in\mathbb{N}$ such that $\varphi_i(\tau_i^m)[S]>0$ for any $m\geq M$. For any $E\in\mathscr{B}(\mathcal{T}_{-i}\times\Theta)$ and $m\geq M$ let measure $\varphi^m_{i,E}$ be given by $D\mapsto\varphi_i(\tau_i^m)[E\cap D]$, for any $D\in\mathscr{B}(\mathcal{T}_{-i}\times\Theta)$. It is immediate that every $\varphi^m_{i,S}$ is absolutely continuous w.r.t. $\varphi_i(\tau_i)$, and thus, by the Radon-Nikodym Theorem, for any $m\geq M$ there exists the Radon-Nikodym derivative of $\varphi^m_{i,S}$ w.r.t. $\varphi_i(\tau_i)$, which we denote by r^m.[36] Now, let belief η^m be defined as follows: $\eta^m[E] = \left(\varphi^m_{i,S^c}\times\mathrm{marg}_{A_{-i}}\eta\right)[E] + \int_E R^m\mathrm{d}\eta$, where $R^m = r^m\circ\mathrm{Proj}_{\mathcal{T}_{-i}\times\Theta}$, for any event $E\in\mathscr{B}(\mathcal{T}_{-i}\times A_{-i}\times\Theta)$.[37] Then:

(*i*) Since $\lim_{m\to\infty}\varphi^m_{i,S^c}=0$, and $\lim_{m\to\infty} r^m = 1$ almost everywhere w.r.t. $\varphi_i(\tau_i)$, it is easy

[36] That is, r^m is the unique, up to a $\varphi_i(\tau_i)$-null set, Borel measurable function such that $\varphi_i(\tau_i^m)[E] = \int_E r^m\mathrm{d}\varphi_i(\tau_i)$ for any event E.

[37] Product measure $\varphi^m_{i,S^c}\times\mathrm{marg}_{A_{-i}}\eta$ is uniquely well-defined due to the Hahn-Kolmogorov Extension Theorem, and the fact that we are working with σ-finite measures.

to check that for any bounded and continuous $f : \mathcal{T}_{-i} \times A_{-i} \times \Theta \to \mathbb{R}$ we have,

$$\lim_{m\to\infty} \int_{\mathcal{T}_{-i}\times A_{-i}\times\Theta} f \mathrm{d}\eta^m = \lim_{m\to\infty} \left[\left(\varphi^m_{i,S^c} \times \mathrm{marg}_{A_{-i}} \eta \right) + \int_{\mathcal{T}_{-i}\times A_{-i}\times\Theta} f \cdot R^m \mathrm{d}\eta \right]$$
$$= \int_{\mathcal{T}_{-i}\times A_{-i}\times\Theta} f \cdot \left(\lim_{m\to\infty} R^m \right) \mathrm{d}\eta = \int_{\mathcal{T}_{-i}\times A_{-i}\times\Theta} f \mathrm{d}\eta.$$

(ii) For any $m \geq M$ and any $E \in \mathscr{B}(\mathcal{T}_{-i} \times \Theta)$,

$$\varphi_i(\tau_i^m)[E] = \varphi^m_{i,S^c}[E] \cdot \left(\mathrm{marg}_{A_{-i}} \eta \right)[A_{-i}] + \int_E r^m \mathrm{d}\varphi_i(\tau_i)$$
$$= \varphi^m_{i,S^c}[E] + \int_{E\times A_{-i}} R^m \mathrm{d}\eta = \left(\mathrm{marg}_{\mathcal{T}_{-i}\times\Theta} \eta^m \right)[E].$$

Thus, $(\eta^n)_{n\in\mathbb{N}}$ converges to η, and there is some $M \in \mathbb{N}$ such that $\mathrm{marg}_{\mathcal{T}_{-i}\times\Theta}\eta^m = \varphi_i(\tau_i^m)$ for any $m \geq M$. ★

CLAIM 2. *For any belief $\eta \in \Delta(\mathcal{T}_{-i} \times A_{-i} \times \Theta)$ such that $BR_i(\eta) = \{a_i\}$ and any sequence of beliefs $(\eta^n)_{n\in\mathbb{N}}$ converging to η, there is some $N \in \mathbb{N}$ such that $a_i \in BR_i(\eta^m)$, for any $m \geq N$.* Denote, for any $a_i' \in A_i$, $F_i(a_i') = U_i\left(\mathrm{marg}_{A_{-i}\times\Theta}\eta, a_i'\right)$, and $F_{i,n}(a_i') = U_i\left(\mathrm{marg}_{A_{-i}\times\Theta}\eta^n, a_i'\right)$, for any $n \in \mathbb{N}$, where,

$$\begin{array}{rccc} U_i: & \Delta(A_{-i}\times\Theta)\times A_i & \longrightarrow & \mathbb{R} \\ & (\eta, a_i) & \to & \displaystyle\int_{A_{-i}\times\Theta} u_i((a_{-i}; a_i), \theta)\, d\eta \end{array}$$

which is proved in Lemma 1 in Weinstein and Yildiz (2013) to be continuous. Then we have that $\lim_{n\to\infty} F_{i,n} = F_i$ pointwise, and from the fact that $BR_i(\eta) = \{a_i\}$, we know that, for every $a_i' \neq a_i$, there is some $n(a_i') \in \mathbb{N}$ such that $F_{i,m}(a_i) > F_{i,m}(a_i')$, for any $m \geq n(a_i')$. Let $N = \max_{a_i'\neq a_i} n(a_i')$. Obviously, $F_{i,m}(a_i) > F_{i,m}(a_i')$ for any $a_i' \neq a_i$ and any $m \geq N$, and thus, $a_i \in BR_i(\eta^m)$ for any $m \geq N$. ★

CLAIM 3. *For any belief $\eta \in \Delta(\mathcal{T}_{-i} \times A_{-i} \times \Theta)$ such that $\eta[\Theta \times \mathrm{Graph}(\mathrm{ICR}_{-i})] = 1$, any sequence of beliefs $(\eta^n)_{n\in\mathbb{N}}$ converging to η, and any $p < 1$, there is some $N_p \in \mathbb{N}$ such that $\eta^m[\Theta \times \mathrm{Graph}(\mathrm{ICR}_{-i})] > p$, for any $m \geq N_p$.* It simply follows from the definition of limit. We know that $\lim_{n\to\infty} \eta^n[\Theta \times \mathrm{Graph}(\mathrm{ICR}_{-i})] = 1$, and thus, that for any $\epsilon > 0$ there is some $N_\epsilon \in \mathbb{N}$ such that $1 - \eta^m[\Theta \times \mathrm{Graph}(\mathrm{ICR}_{-i})] < \varepsilon$ for any $m \geq M$. Set $\varepsilon = 1 - p$ and $N_p = N_\varepsilon$ to have the proof completed. ★

PROOF OF THE THEOREM. Take standard hierarchy τ_i, sequence $(\tau_i^n)_{n\in\mathbb{N}}$ converging to τ_i, and $a_i \in \mathrm{ICSR}_i(\tau_i)$. Then, there is some $\eta \in \Delta(\mathcal{T}_{-i} \times A_{-i} \times \Theta)$ such that: (i) $\mathrm{marg}_{\mathcal{T}_{-i}\times\Theta}\eta = \varphi_i(\tau_i)$, (ii) $BR_i(\eta) = \{a_i\}$, and (iii) $\eta[\Theta \times \mathrm{Graph}(\mathrm{ICR}_{-i})] = 1$. Take sequence $(\eta^n)_{n\in\mathbb{N}}$ converging to η as in Claim 1, and fix $p < 1$. We know from Claims

1, 2 and 3 that there is some $n_p = \max\{M, N, N_p\}$ such that $a_i \in BR_i(\eta^m)$ and $\eta^m[\Theta \times \text{Graph}(\text{ICR}_{-i})] > p$ for any $m \geq p$. Thus, there is some $n_p \in \mathbb{N}$ such that $a_i \in \text{ICR}_i^p(\tau_i^m)$ for any $m \geq n_p$. ■

Proof. Fix $i \in I$ and $\tau_i \in \mathcal{T}$ such that $|\text{ICSR}_i(\tau_i)| > 1$. Then, we know from Theorem 2 that, for any $p < 1$ and any $(\tau_i^n)_{n\in\mathbb{N}}$ converging to τ_i, there is some $n_p \in \mathbb{N}$ such that $|\text{ICR}_i^p(\tau_i^m)| > 1$, for any $m \geq n_p$, and thus, that no sequence in $\mathcal{U}_i^p$ converges to τ_i. Hence, τ_i is not in the closure of $\mathcal{U}_i^p$. ■

REFERENCES

Aliprantis, Charambolos D. and Kim C. Border (1999). *Infinite dimensional analysis: a hitchhiker's guide, 3rd edition*. Springer, Berlin.

Armbruster, Walter and Werner Böge (1979). Bayesian game theory. In: *Game theory and related topics*. North Holland, Amsterdam.

Battigalli, Pierpaolo, Alfredo Di Tillio, Eduardo Grillo and Antonio Penta (2011). "Interactive epistemology and solution concepts for games with asymmetric information". *The B.E. Journal of Theoretical Economics* **11**, Article 6.

Battigalli, Pierpaolo and Marciano Siniscalchi (2003). "Rationalization and incomplete information". *The B.E. Journal of Theoretical Economics* **3**, 1–46.

Battigalli, Pierpaolo and Marciano Siniscalchi (2007). "Interactive epistemology in games with payoff uncertainty". *Research in Economics* **61**, 165–184.

Bergemann, Dirk and Stephen Morris (2005). "Robust mechanism design". *Econometrica* **73**, 1771–1813.

Bergemann, Dirk and Stephen Morris (2013). "Robust predictions in games with incomplete information". *Econometrica* **81**, 1251–1308.

Bernheim, B. Douglas (1984). "Rationalizable strategic behaviour". *Econometrica* **52**, 1007–1028.

Böge, Werner and Th. Eisele (1979). "On solutions of Bayesian games". *The International Journal of Game Theory* **8**, 193–215.

Brandenburger, Adam and Eddie Dekel (1987). "Rationalizability and correlated equilibria". *Econometrica* **55**, 1391–1402.

Brandenburger, Adam and Eddie Dekel (1993). "Hierarchies of beliefs and common knowledge". *Journal of Economic Theory* **59**, 189–198.

Chen, Yi-Chun (2012). "A structure theorem for rationalizability in the normal form of dynamic games". *Games and Economic Behavior* **75**, 587–597.

Chen, Yi-Chun, Alfredo Di Tillio, Eduardo Faingold and Siyang Xiong (2010). "Uniform topologies on types". *Theoretical Economics* **5**, 445–478.

Dekel, Eddie and Marciano Siniscalchi (2013). "Epistemic game theory". Mimeo.

Dekel, Eddie, Drew Fudenberg and Stephen Morris (2006). "Topologies on types". *Theoretical Economics* **1**, 275–309.

Dekel, Eddie, Drew Fudenberg and Stephen Morris (2007). "Interim correlated rationalizability". *Theoretical Economics* **2**, 15–40.

Ely, Jeffrey C. and Marcin Pęski (2006). "Hierarchies of beliefs and interim rationalizability". *Theoretical Economics* **1**, 19–65.

Germano, Fabrizio and Peio Zuazo-Garin (2015). "Bounded rationality and correlated equilibria". Mimeo.

Harsanyi, John C. (1967–1968). "Games with incomplete information played by 'Bayesian' players, I–III". *Management Science* **14**, 159–182, 320–334, 486–502.

Heifetz, Aviad and Dov Samet (1998). "Topology-free typology of beliefs". *Journal of Economic Theory* **82**, 324–341.

Hu, Tai-Wei (2007). "On p-rationalizability and approximate common certainty of rationality". *Journal of Economic Theory* **136**, 379–391.

Kajii, Atsushi and Stephen Morris (1997). "The robustness of equequilibria to incomplete infomration". *Econometrica* **65**, 1283–1309.

Mertens, Jean-François and Shmuel Zamir (1985). "Formulation of Bayesian analysis for games with incomplete information". *International Journal of Game Theory* **14**(1), 1–29.

Monderer, Dov and Dov Samet (1989). "Approximating common knowledge with common beliefs". *Games and Economic Behavior* **1**, 170–190.

Osborne, Martin J. and Ariel Rubinstein (1994). *A course in game theory*. MIT Press Books.

Pearce, David G. (1984). "Rationalizable strategic behavior and the problem of perfection". *Econometrica* **52**, 1029–1050.

Penta, Antonio (2012). "Higher order uncertainty and information: static and dynamic games". *Econometrica* **80**, 631–660.

Penta, Antonio (2013). "On the structure of rationalizability for arbitrary spaces of uncertainty". *Theoretical Economics* **8**, 405–430.

Piermont, Evan and Peio Zuazo-Garin (2015). "From common knowledge to common belief". Mimeo.

Rubinstein, Ariel (1989). "The electronic mail game: strategic behaviour under 'almost common knowledge"'. *American Economic Review* **79**, 385–391.

Tan, Tommy Chin-Chiu and Sérgio Ribeiro da Costa Werlang (1988). "The Bayesian foundation of solution concepts of games". *Journal of Economic Theory* **45**, 370–391.

Weinstein, Jonathan and Muhamet Yildiz (2007). "A structure theorem for rationalizability with application to robust predictions of refinements". *Econometrica* **75**, 365–400.

Weinstein, Jonathan and Muhamet Yildiz (2013). "Properties of interim correlated rationalizability in infinite games". Mimeo.