

Uncertain Rationality, Depth of Reasoning and Robustness in Games with Incomplete Information*

Fabrizio Germano[†] Jonathan Weinstein[‡] Peio Zuazo-Garin[§]

July 2, 2017

ABSTRACT

Predictions under common knowledge of payoffs may differ from those under arbitrarily, but finitely, many orders of mutual knowledge; Rubinstein's (1989) *Email game* is a seminal example. Weinstein and Yildiz (2007) showed that the discontinuity in the example generalizes: for all types with multiple rationalizable (ICR) actions, there exist similar types with unique rationalizable action. This paper studies how a wide class of departures from common belief in rationality impact Weinstein and Yildiz's discontinuity. We weaken ICR to ICR^λ , where λ is a sequence whose n^{th} term is the probability players attach to $(n - 1)^{\text{th}}$ -order belief in rationality. We find that Weinstein and Yildiz's discontinuity holds when higher-order belief in rationality remains above some threshold (constant λ), but fails when higher-order belief in rationality eventually becomes low enough (λ converging to 0).

*The authors thank the helpful comments and discussion by Pierpaolo Battigalli, Eddie Dekel, Aviad Heifetz, Tai-Wei Hu, Willemien Kets, Stephen Morris, Antonio Penta and seminar and conference audiences at Universitat de Barcelona, Washington University in St. Louis, University of Wisconsin-Madison, the NASM'16, LOFT11 and GAMES 2016. Germano acknowledges financial support from the Spanish Ministry of Economy and Competitiveness (Grant ECO2014-59225-P), 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 Competitiveness and from the Department of Language Policy and Culture of the Basque Government (grants ECO2012-31346 and POS-2015-1-0022 and IT568-13, respectively), and thanks Northwestern University (both the Department of Economics and the Math Center at MEDS Department) 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.

[‡]Washington University in St. Louis, Department of Economics; j.weinstein@wustl.edu.

[§]University of the Basque Country (UPV/EHU) and *BRiDGE* Group; Department of Foundations of Economic Analysis I, Avenida Lehendakari Agirre 83, E-48015, Bilbao, Spain; peio.zuazo@ehu.eus.

KEYWORDS. Robustness, rationalizability, bounded rationality, incomplete information, belief hierarchies.

JEL CLASSIFICATION. C72, D82, D83.

1 INTRODUCTION

An extensive literature has taught us that small perturbations to players' beliefs may induce large changes in our strategic predictions. In particular, Rubinstein's email game (detailed below) is a seminal example along these lines: it showed that predictions under common knowledge of payoffs may differ from those under arbitrarily, but finitely, many orders of mutual knowledge. That is, if I know that you know that I know, etc., what the payoffs are, but this chain breaks after finitely many levels, some outcomes which would be rationalizable under full common knowledge will be non-rationalizable under this partial knowledge. [Weinstein and Yildiz \(2007\)](#) showed that the discontinuity in the example generalizes: for any type with multiple rationalizable actions, there are types with very similar beliefs which have a unique rationalizable action. We call this the "WY-discontinuity." The notion of "small" changes in beliefs, i.e. the topology on types, is of course highly relevant here: we use here the product topology, as in [Weinstein and Yildiz \(2007\)](#). The significance of this choice is that arbitrary changes in very high-order beliefs are measured as small. Alternatively, recent papers such as [Chen, Luo and Qu \(2015\)](#) have shown that requiring *uniform* convergence of belief hierarchies does imply convergence of strategic behavior.

The main result of [Weinstein and Yildiz \(2007\)](#) uses the solution concept of interim correlated rationalizability (ICR), which is founded on the assumption of common belief in rationality. In this paper we ask: what happens to the WY-discontinuity if we weaken this assumption? Specifically, we weaken ICR to the more permissive *interim correlated λ -rationalizability* (ICR^λ) where $\lambda = (\lambda_n)_{n \in \mathbb{N}}$ is a sequence of probabilities with the interpretation that λ_n is the reliability that players attach to n^{th} -order belief in rationality; ICR itself would be the special case that $\lambda = (1, 1, \dots)$. The answer is twofold: when λ is constant in n and close enough to 1 we find that WY-discontinuity remains (Proposition 4), but when $(\lambda_n)_{n \in \mathbb{N}} \rightarrow 0$ as $n \rightarrow \infty$ we find that continuity is restored (Proposition 1). That is, when common belief in rationality breaks down almost completely at high orders, the continuity of behavior with respect to perturbations of belief hierarchies is restored. As we discuss in Section 3.2, the ICR^λ concept is very flexible; as λ varies it covers concepts close

to ICR as well as those much further away (such as rationality without any mutual belief in rationality.)

This restoration of continuity is important, because, as discussed in [Weinstein and Yildiz \(2007\)](#), the WY-discontinuity has profound implications for the large applied literature on equilibrium refinements. When the discontinuity obtains, all non-trivial refinements are non-robust to the introduction of incomplete information, or to changes in the assumptions on players' information. Here we show that some (but not all) relaxations of common knowledge of rationality restore continuity and hence the possibility of robust refinements.

In addition to our main results in Propositions 4 and 1 we also prove some standard robustness properties of ICR^λ . We show that different types that induce the same belief hierarchy induce the same set of ICR^λ actions (type-representation invariance, Proposition 1). We also show that, for each fixed λ , ICR^λ is an upper-hemicontinuous correspondence, that is, small misspecifications of beliefs do not give rise to unexpected behavior (Proposition 2). Regarding robustness to the weakening of common belief in rationality, we show that, when the belief hierarchy is fixed, correspondence ICR^λ , varying on λ , is upper-hemicontinuous everywhere and is lower-hemicontinuous at $\lambda = (1, 1, \dots, 1, \dots)$, where it coincides with ICR (Proposition 3). This result establishes the full robustness of ICR to a slight weakening of common belief in rationality. Finally, we provide an epistemic foundation of ICR^λ to show that it characterizes rationality and common λ -belief in rationality, thus confirming its suitability for the formalization of perturbations in common belief in rationality (Proposition 5). In particular, all these results, besides Theorems 4 and 1, are formulated for generic λ and are therefore applicable to a variety of well-known solution concepts obtained by considering particular subfamilies of λ (e.g. ICR, p -rationalizability or k -level rationalizability).

1.1 RUBINSTEIN'S EMAIL GAME

The incomplete information game given by the following payoff matrix is an adaptation of Rubinstein's game:

	Attack	No attack	
Attack	$\theta \quad \theta$	$\theta - 1 \quad 0$	for $\theta \in \Theta = \{-2/5, 2/5\}$.
No attack	$0 \quad \theta - 1$	$0 \quad 0$	

Ex ante, players assign probability $1/2$ to each of the values $-2/5$ and $2/5$. Player 1 observes the value of θ and automatically sends a message to Player 2, if $\theta = 2/5$. Each player automatically sends a message back whenever he receives one, and each message is lost, with probability $1/2$. When a message is lost, the process automatically stops and each player takes one of the actions *Attack* or *No attack*. This game can be modeled by the type space $T = \{-1, 1, 3, 5, \dots\} \times \{0, 2, 4, 6, \dots\}$, where the type t_i is the total number of messages sent or received by player i (except for type $t_1 = -1$, who knows that $\theta = -2/5$), and the common prior μ on $T \times \Theta$, where $\mu(\theta = -2/5, t_1 = -1, t_2 = 0) = 1/2$ and for each integer $m \geq 1$, $\mu(\theta = 2/5, t_1 = 2m-1, t_2 = 2m-2) = 1/2^{2m}$ and $\mu(\theta = 2/5, t_1 = 2m-1, t_2 = 2m) = 1/2^{2m+1}$. Here, for $k \geq 1$, type k knows that $\theta = 2/5$, knows that the other player knows $\theta = 2/5$, and so on, through k orders. Now, type $t_1 = -1$ knows that $\theta = -2/5$ and, hence, his unique rationalizable action is *No attack*. Type $t_2 = 0$ does not know θ but puts probability $2/3$ on type $t_1 = -1$, thus believing that player 1 will play *No attack* with at least probability $2/3$, so that *No attack* is the only best reply and, hence, the only rationalizable action. Applying this argument inductively for each type k , one concludes that the new incomplete-information game is dominance-solvable and the unique rationalizable action for all types is *No attack*.

Consider Rubinstein's commentary on his example: "It is hard to imagine that [when many messages are sent] a player will not play [according to the Pareto-dominant equilibrium.] The sharp contrast between our intuition and the game-theoretic analysis is what makes this example paradoxical. The example joins a long list of games...in which it seems that the source of the discrepancy is rooted in the fact that in our formal analysis we use mathematical induction while human beings do not use mathematical induction when reasoning. Systematic explanation of our intuition...is definitely a most intriguing question." Indeed, the goal of this paper is to formalize this intuition. Our main results will show that some weakenings of the inductive reasoning of rationalizability will maintain the unique, counterintuitive selection in the email example which underlies the WY-discontinuity, while others will return us to the more intuitive case of multiple equilibria, from which we may select according to a criterion such as Pareto-dominance.

When reasoning is the same at every level, even if it assigns less than full confidence to opponents' rationality, the unique selection persists. Indeed, assume that each player i assigns probability $p > 2/3$ to the other player being rational, assigns probability p to the other player assigning probability p to i being rational, and so on. Again, type $t_1 = -1$ knows that $\theta = -2/5$ and, hence, plays *No attack*, regardless

of her beliefs about the other player's choice. Type $t_2 = 0$ does not know θ but puts probability $2/3$ on type $t_1 = -1$, thus believing that player 1 will play *No attack* with at least probability $p \cdot 2/3 > 2/5$, so that *No attack* is the only best reply, and hence, the only p -rationalizable action. Similarly, type $t_1 = 1$ puts probability $2/3$ on type $t_2 = 0$, and thus will play *No attack* with at least probability $p \cdot 2/3 > 2/5$, so that, again, *No attack* is the only best reply, and, hence, the only p -rationalizable action, and so on. This is an example of Proposition 4: under appropriate conditions, there will always be a $p < 1$ large enough that unique selection survives in this way.

The opposite result obtains when players lose almost all confidence in their reasoning at later iterations. Specifically, assume that each player i assigns probability λ_1 to the other player being rational, assigns probability λ_2 to the other player assigning probability λ_1 to i being rational, and so on, where $\lambda_k \rightarrow 0$, so that the effect of higher-order restrictions vanishes as we move up in the hierarchy. Now note that even if $t_2 = k - 1$ is a type that always plays *No attack*, if $\lambda_k < 2/5$, we cannot guarantee that *No attack* is the only best reply for type $t_1 = k$. Thus, we can always find a sufficiently high number of messages for which the action *Attack* survives the iterated deletion procedure. This is an example of Proposition 1: when confidence in higher-order reasoning breaks down at high orders, all strictly rationalizable actions will be rationalizable in any perturbation.

1.2 OTHER RELATED LITERATURE

This paper scrutinizes the discontinuity in the rationalizable set by altering the solution concept. Specifically, it studies the impact of weakening common belief in rationality on the WY-discontinuity, in the spirit of the quote above from Rubinstein (1989). Also in this line, previous papers have studied the effects of departure from the standard rationality benchmark by invoking finite depth of reasoning assumptions. Strzalecki (2010) and Heifetz and Kets (2016) extend the notion of type/belief hierarchy so that it incorporates uncertainty and higher-order beliefs about the depth of reasoning. Within this richer framework Heifetz and Kets (2016) perturb common belief in infinite depth of reasoning (an implicit feature of the standard notion of type in Weinstein and Yildiz (2007)) and find that under *almost* common belief in infinite depth of reasoning, the corresponding notion of ICR does not exhibit the WY-discontinuity.¹

A second research agenda spawned by the finding of discontinuities in rational-

¹The connection of this paper and Heifetz and Kets (2016) is examined in further detail in Section 4.2.2.

izability considered replacing the product topology with alternate notions of proximity. [Dekel, Fudenberg and Morris \(2006\)](#) introduce the *strategic topology* which is implicitly defined as the coarsest topology for the space of belief hierarchies under which ICR is upper-hemicontinuous and strict ICR is lower-hemicontinuous. Previous papers by [Monderer and Samet \(1996\)](#) and [Kajii and Morris \(1997\)](#) ensure the robustness of equilibria under incomplete information by proposing topologies whose corresponding notion of perturbation, based on common p -belief, require (unlike perturbations in the product topology) approximations to take similarity of all higher-order beliefs into account. Recent work by [Chen, Di Tillio, Faingold and Xiong \(2010\)](#); [Chen, Tillio, Faingold and Xiong \(2016\)](#) bridges the gap between the two approaches by providing the exact metric that characterizes the strategic topology and some of its refinements.

Finally, in a third category, an important branch of the literature exploits discontinuities of behavior to construct equilibrium selection arguments (e.g., [Carlsson and van Damme \(1993\)](#)), explain large changes on behavior induces by small changes in economic fundamentals (e.g., [Morris and Shin \(1998\)](#)), and extend the domain in which the WY-discontinuity holds to dynamic games ([Penta \(2012\)](#) and [Chen \(2012\)](#)) and to more general cases of payoff uncertainty ([Penta \(2013\)](#), [Chen, Takahashi and Siyan \(2014\)](#); [Chen, Takahashi and Xiong \(2014\)](#)).

2 PRELIMINARIES

In this section we briefly review some well-known ideas central to our study. First, in Section 2.1 we describe the game-theoretical framework employed to model interaction. This will consist of games with incomplete information and Bayesian games. Remember that in such games the uncertainty each player faces is twofold: it refers to states of nature that affect preferences (*payoff uncertainty*) and to the actions the rest of players choose (*strategic uncertainty*). Payoff uncertainty is dealt with by exogenously setting either types as defined by [Harsanyi \(1967–1968\)](#) or belief hierarchies. The construction of the latter, together with that of universal type space, is recalled in Section 2.2. Strategic uncertainty is endogenously resolved by means of a solution concept, namely interim correlated rationalizability. This is presented in Section 2.3, where we also recall the Structure Theorem of [Weinstein and Yildiz \(2007\)](#) and some of its implications.

2.1 GAMES WITH INCOMPLETE INFORMATION AND BAYESIAN GAMES

A (static) game with incomplete information consists of a list $\mathcal{G} = \langle I, \Theta, (A_i, u_i)_{i \in I} \rangle$, where: (i) I is a finite set of players, (ii) Θ is a compact, metrizable set of *payoff states*, and for each player i we have (iii) a finite set of *actions*, A_i , and (iv) a continuous *utility function* $u_i : A \times \Theta \rightarrow \mathbb{R}$, where $A = \prod_{i \in I} A_i$ is the set of *action profiles*. For each player i , we refer to a probability measure $\mu_i \in \Delta(A_{-i} \times \Theta)$,² where $A_{-i} = \prod_{j \neq i} A_j$, as a *conjecture*, and we define player i 's *best-reply correspondence* as,

$$\begin{aligned} BR_i : \Delta(A_{-i} \times \Theta) &\rightrightarrows A_i \\ \mu_i &\mapsto \arg \max_{a_i \in A_i} \int_{A_{-i} \times \Theta} u_i((a_{-i}; a_i), \theta) d\mu_i, \end{aligned}$$

which, due to the topological assumptions specified above, is known to be both non-empty and upper-hemicontinuous.³ We typically represent players' beliefs over Θ by endowing \mathcal{G} with a *type structure* à la Harsanyi (1967–1968). A type structure is a list $\mathcal{T} = \langle T_i, \pi_i \rangle_{i \in I}$ where for each player i we have: (i) a compact and metrizable set of *types*, T_i , and (ii) a continuous *belief map* $\pi_i : T_i \rightarrow \Delta(T_{-i} \times \Theta)$ where $T_{-i} = \prod_{j \neq i} T_j$. We refer to a pair $\langle \mathcal{G}, \mathcal{T} \rangle$ as a *Bayesian game*. The beliefs in type structures are not assumed to arise from a common prior.

2.2 BELIEF HIERARCHIES AND UNIVERSAL TYPE SPACE

We follow Brandenburger and Dekel's (1993) formulation of universal type space. For each player i set first $X_i^1 = \Theta$ and $Z_i^1 = \Delta(X_i^1)$, and call each element $\tau_{i,1} \in Z_i^1$ *first-order belief*. Then, set recursively $X_i^{n+1} = X_i^n \times \prod_{j \neq i} Z_j^n$ and $Z_i^{n+1} = \Delta(X_i^n)$ for any $n \in \mathbb{N}$. We refer to each $\tau_{i,n} \in Z_i^n$ as n^{th} -order belief, and to the elements of $\mathcal{T}_i^0 = \prod_{n \in \mathbb{N}} Z_i^n$, as *belief hierarchies*. Belief hierarchy τ_i is said to be *coherent* if higher-order belief do not contradict lower order ones, i.e., if $\text{marg}_{X_i^n} \tau_{i,n+1} = \tau_{i,n}$ for any $n \in \mathbb{N}$. Let \mathcal{T}_i^1 denote the set of coherent belief hierarchies and \mathcal{T}_i , the set of belief hierarchies that exhibit common belief in coherence.⁴ Brandenburger

²For 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, $\mathcal{B}(X)$. Topologies for the other spaces are standard: the induced topology for subsets and the product topology for Cartesian products.

³When necessary, with some abuse of notation we will write $BR_i(\eta_i) = BR_i(\text{marg}_{A_{-i} \times \Theta} \eta_i)$ for any compact and metrizable space X and any belief $\eta_i \in \Delta(X \times A_{-i} \times \Theta)$.

⁴Formally, $\mathcal{T}_i = \bigcap_{n \geq 0} \mathcal{T}_i^n$, where $\mathcal{T}_i^{n+1} = \{\tau_i \in \mathcal{T}_i^n \mid \tau_{i,m}[\text{Proj}_{X_i^m}(\mathcal{T}_i^n \times \Theta)] = 1 \text{ for any } m \in \mathbb{N}\}$ for each $n \in \mathbb{N}$, being $\mathcal{T}_i^n = \prod_{j \neq i} \mathcal{T}_j^n$. For any product space $X \times Y$ and any subset $S \subseteq X \times Y$, we denote *projections* on some component of X by $\text{Proj}_X S = \{x \in X \mid (x, y) \in S \text{ for some } y \in Y\}$.

and Dekel (1993) show that there exists a homeomorphism $\varphi_i : \mathcal{T}_i \rightarrow \Delta(\mathcal{T}_{-i} \times \Theta)$, with $\mathcal{T}_{-i} = \prod_{j \neq i} \mathcal{T}_j$, such that $\text{marg}_{X_i^n} \varphi_i(\tau_i) = \tau_{i,n}$ for any belief hierarchy τ_i and any $n \in \mathbb{N}$. Obviously, $\langle \mathcal{T}_i, \varphi_i \rangle_{i \in I}$ is a type structure for game with incomplete information \mathcal{G} ; we refer to it as the *universal type space*.

Throughout the above constructions, as is standard we topologize spaces of beliefs by the weak* topology and product spaces by the product topology, and in this way the space of belief hierarchies inherits a topology. A corresponding metric is also inherited at each step of the recursion: first normalize the metric on the basic space Θ so its diameter is at most 1 (this property will be inherited at each step.) Then apply the Prohorov metric to spaces of beliefs, the sup metric to finite products, and the discounted metric,

$$d(x, x') = \sum_{n=1}^{\infty} 2^{-n} d(x_n, x'_n)$$

to infinite product spaces. Thus the space of belief hierarchies also inherits a metric structure.

Finally, it is known that for each type structure \mathcal{T} , each type t_i induces a belief hierarchy $\tau_i(t_i) = (\tau_{i,n}(t_i))_{n \in \mathbb{N}}$ as follows: consider first-order belief $\tau_{i,1}(t_i) = \text{marg}_{\Theta} \pi_i(t_i)$ and then, for any $n \in \mathbb{N}$ define $(n+1)^{\text{th}}$ -order belief $\tau_{i,n+1}(t_i)$ by setting,

$$\tau_{i,n+1}(t_i)[E_{n+1}] = \pi_i(t_i) [\{(t_{-i}, \theta) \in T_{-i} \times \Theta \mid (\tau_{-i,n}(t_{-i}), \theta) \in E_{n+1}\}],$$

for any measurable $E_{n+1} \subseteq \mathcal{T}_{-i}^{n+1} \times \Theta$. The recursive construction being well-defined follows from the fact that, as proved by Brandenburger and Dekel (1993), every $\tau_{i,n} : T_i \rightarrow Z_i^n$ is continuous. In addition, it is easy to see that $\tau_i(T_i) \subseteq \mathcal{T}_i$; thus, $\tau_i : T_i \rightarrow \mathcal{T}_i$ is a well-defined continuous map. Furthermore, if T_i has non-redundant types,⁵ then it is homeomorphic to $\tau_i(T_i)$.

2.3 RATIONALIZABILITY AND THE WY-DISCONTINUITY

Once a player's uncertainty w.r.t. the set of payoff states is formalized by means of some type or belief hierarchy, it becomes pertinent to wonder which subset of actions constitutes a *reasonable* choice at the interim stage. By "reasonable" we will mean those actions consistent with rationality and common belief in opponents' rationality, or, in other words, to those actions that survive iterated deletion of strictly dominated actions. This idea is formalized by interim correlated rationalizability (ICR),

⁵That is, if every two distinct types induce different belief hierarchies: $t_i \neq t'_i$ implies that $\tau_i(t_i) \neq \tau_i(t'_i)$.

introduced by [Dekel, Fudenberg and Morris \(2007\)](#). Next we recall the version of the definition of ICR due to [Battigalli, Di Tillio, Grillo and Penta \(2011\)](#). Given a game with incomplete information, \mathcal{G} , player i 's set of (*interim correlated*) *rationalizable* (ICR) actions for belief hierarchy τ_i is defined as $\text{ICR}_i(\tau_i) = \bigcap_{n \geq 0} \text{ICR}_{i,n}(\tau_i)$, where $\text{ICR}_{i,0}(\tau_i) = A_i$ and, recursively,⁶

$$\text{ICR}_{i,n}(\tau_i) = \left. \begin{array}{l} a_i \in A_i \\ \left. \begin{array}{l} \text{There exists some measurable } \sigma_{-i} : \mathcal{T}_{-i} \times \Theta \rightarrow \Delta(A_{-i}) \text{ such that:} \\ (i) \quad \text{supp } \sigma_{-i}(\tau_{-i}, \theta) \subseteq \text{ICR}_{-i,n-1}(\tau_{-i}) \text{ for } \varphi_i(\tau_i)\text{-a.e. } (\tau_{-i}, \theta) \in \mathcal{T}_{-i} \times \Theta, \\ (ii) \quad a_i \in \arg \max_{a'_i \in A_i} \int_{\mathcal{T}_{-i} \times \Theta} \left(\sum_{a_{-i} \in A_{-i}} \sigma_{-i}(\tau_{-i}, \theta)[a_{-i}] \cdot u_i((a_{-i}, a'_i), \theta) \right) d\varphi_i(\tau_i) \end{array} \right\} \end{array} \right\},$$

for any $n \in \mathbb{N}$. A similar definition can be given when the original input is a Bayesian game $\langle \mathcal{G}, \mathcal{T} \rangle$ instead of a game with incomplete information: making the obvious changes player i 's set of rationalizable actions for type t_i , $\text{ICR}_i^{\mathcal{T}}(t_i)$, is obtained.⁷ [Dekel, Fudenberg and Morris \(2007\)](#) and [Battigalli, Di Tillio, Grillo and Penta \(2011\)](#) show that the two definitions are consistent: the set of rationalizable actions corresponding to a type coincides with the set of rationalizable actions corresponding to the belief hierarchy induced by the type.⁸ We refer to this property of ICR as *type-representation invariance*. In addition, it is shown by [Dekel, Fudenberg and Morris \(2007\)](#) that correspondence $\text{ICR}_i : \mathcal{T}_i \rightrightarrows A_i$ is upper-hemicontinuous, and by [Dekel, Fudenberg and Morris \(2007\)](#) and [Battigalli, Di Tillio, Grillo and Penta \(2011\)](#), that rationalizability characterizes behavior under rationality and common belief in rationality.

In their the study of ICR, [Weinstein and Yildiz \(2007\)](#) find a striking property that generalizes the discontinuity in the Email game from an isolated phenomenon to a general feature of games with incomplete information. If \mathcal{G} is such that Θ satisfies the *richness* condition, i.e., that it is not commonly known that some action is not strictly dominant,⁹ then for any belief hierarchy τ_i and any $a_i \in \text{ICR}_i(\tau_i)$ there exists

⁶In addition, let us denote $\text{ICR}_{-i,n-1}(\tau_{-i}) = \prod_{j \neq i} \text{ICR}_{j,n-1}(\tau_j)$.

⁷Simply replace, in the previous definition, belief hierarchies for types and the homeomorphism corresponding to the universal type space for the belief map corresponding to the type structure.

⁸That is, for any Bayesian game $\langle \mathcal{G}, \mathcal{T} \rangle$, any player i and any type t_i it holds that $\text{ICR}_i^{\mathcal{T}}(t_i) = \text{ICR}_i(\tau_i(t_i))$.

⁹That is, for any player i and any action a_i there exists some payoff state θ_{a_i} that makes a_i strictly dominant for player i .

some sequence $(\tau_i^n)_{n \in \mathbb{N}}$ converging to τ_i such that $\text{ICR}_i(\tau_i^n) = \{a_i\}$ for any $n \in \mathbb{N}$.¹⁰ This property, which we refer to as the *WY-discontinuity*, has important implications for games with incomplete information:

- *Non-robustness of refinements.* No non-trivial refinement of ICR is robust in the sense of upper-hemicontinuity on \mathcal{T}_i . To see why, suppose that $S_i : \mathcal{T}_i \rightrightarrows A_i$ is a non-trivial refinement of ICR_i . Then, there exists some belief hierarchy τ_i such that $\text{ICR}_i(\tau_i) \setminus S_i(\tau_i)$ contains some action a_i . By [Weinstein and Yildiz's \(2007\)](#) result, we know that there exists some sequence $(\tau_i^n)_{n \in \mathbb{N}}$ such that $\emptyset \neq S_i(\tau_i^n) \subseteq \text{ICR}_i(\tau_i^n) = \{a_i\}$ for any $n \in \mathbb{N}$; hence S_i cannot be upper-hemicontinuous. In particular, the fact that equilibrium outcomes refine ICR outcomes implies that equilibrium predictions are not robust: small misspecifications of players' uncertainty by the analyst lead to outcomes overlooked in the original model.
- *Generic uniqueness of rationalizability.* There exists an open and dense subset of \mathcal{T}_i such that the set of ICR_i actions corresponding to each belief hierarchy in the set is unique. Thus, rationalizability generically (in a particular topological sense) yields a unique prediction.

3 INTERIM CORRELATED λ -RATIONALIZABILITY

3.1 DEFINITION

We now introduce interim correlated λ -rationalizability (ICR^λ), the solution concept that will formalize our relaxation of common belief in rationality. This concept captures the ideas that (A) rationality may not be common belief and (B) players' confidence in the rationality of others may be different at different orders. The sequence $\lambda \in [0, 1]^{\mathbb{N}}$ signifies that when reasoning at order n , players have confidence λ_n in the rationality of others, as captured in the following definition:

DEFINITION 1 (Interim correlated λ -rationalizability). *Let \mathcal{G} be a game with incomplete information and λ , a sequence of probabilities. Then, player i 's set of (interim correlated) λ -rationalizable actions for belief hierarchy τ_i is defined as $\text{ICR}_i^\lambda(\tau_i) = \bigcap_{n \geq 0} \text{ICR}_{i,n}^\lambda(\tau_i)$, where $\text{ICR}_{i,0}^\lambda(\tau_i) = A_i$ and $C_{i,0}^\lambda(\tau_i) = \{\eta_i \in \Delta(\mathcal{T}_{-i} \times A_{-i} \times \Theta) \mid \text{marg}_{\mathcal{T}_{-i} \times \Theta} \eta_i =$*

¹⁰Recently, [Penta \(2013\)](#) found that the rather demanding richness condition can be abandoned and the discontinuity result extended to relatively mild relaxations of common knowledge assumptions.

$\varphi_i(\tau_i)\}$, and recursively, for any $n \in \mathbb{N}$,

$$\begin{aligned} \text{ICR}_{i,n}^\lambda(\tau_i) &= \{a_i \in \text{ICR}_{i,n-1}^\lambda(\tau_i) \mid a_i \in BR_i(\eta_i) \text{ for some } \eta_i \in C_{i,n-1}^\lambda(\tau_i)\}, \\ C_{i,n}^\lambda(\tau_i) &= \{\eta_i \in C_{i,n-1}^\lambda(\tau_i) \mid \eta_i[M] \geq \lambda_n \text{ for some measurable } M \subseteq \text{Graph}(\text{ICR}_{-i,n}^\lambda) \times \Theta\}. \end{aligned}$$

We provide below in Remark 1 a more natural characterization of ICR^λ . For $p \in [0, 1]$, we will use $\lambda = \bar{p}$ to signify the constant sequence $\lambda_n \equiv p$. Then $\text{ICR}^{\bar{p}}$ will reflect reasoning that is depth-independent, capturing departures from common belief in rationality in the sense of (A), but not (B) above. The case of decreasing λ represents depth-dependent reasoning, where we reason less confidently at higher orders, hence capturing both (A) and (B) above. We will especially consider the case $\lambda \rightarrow 0$, which represents a near-complete breakdown in confidence of others' reasoning at high orders.

3.2 SPECIAL CASES OF ICR^λ

Let $\Lambda = [0, 1]^{\mathbb{N}}$ represent the set of probability sequences. Certain subsets of Λ give rise to different well-known solutions concepts as special cases of ICR^λ :

- (i) *p-Rationalizability.* $\lambda = \bar{p}$, for any $p \in [0, 1]$. These sequences follow the idea by [Monderer and Samet \(1987\)](#) of perturbing common belief by employing p -beliefs; this approach was also followed by [Hu \(2007\)](#) in his analysis of robustness to perturbation in common belief in rationality in the context of games with complete information. We sometimes refer to $\text{ICR}^{\bar{p}}$ actions as *interim correlated p-rationalizable*.
- (ii) *Rationalizability.* The special case $\lambda = \bar{1}$. The standard case of common belief in rationality, that is, infinite depth of reasoning in which player adhere probability 1 to rationality at every iteration. The case $\text{ICR}^{\bar{1}}$ reduces to the standard notion of ICR, as defined by [Dekel, Fudenberg and Morris \(2007\)](#) and discussed above.
- (iii) *Models with k orders of belief in rationality.* For each $k \geq 0$ define sequence $\lambda^k = (\lambda_n^k)_{n \in \mathbb{N}}$ by

$$\lambda_n^k = 1 \text{ if } n \leq k \text{ and } \lambda_n^k = 0 \text{ otherwise.}$$

An action $a_i \in \text{ICR}_i^{\lambda^k}(\tau_i)$ corresponds to the choice of a player who assumes

that others are rational for $k - 1$ orders and makes no further assumptions.¹¹

(iv) *Models with distinct “cognitive bound” and “rationality bound”*. [Friedenberg, Kets and Kneeland \(2016\)](#) define the following (on p. 3):

- Rationality: Say Ann is *rational* if she maximizes her expected utility given subjective belief about how Bob plays the game.
- Cognition: Say Ann is *cognitive* if she has a subjective belief about how Bob plays the game.

From this they further define:

- Reasoning About Rationality: Say that Ann has a *rationality bound* of level n if she is rational, thinks that Bob is rational, thinks that Bob thinks she is rational, and so on up to the statement that includes the word “rational” n times, but no further.
- Reasoning About Cognition: Say that Ann has a *cognitive bound* of level m if she is thinking about Bob’s strategy choice, if she is thinking about what Bob is thinking about her strategy choice, and so on up to the statement that includes the word “thinking” m times, but no further.

Since rationality is stronger than cognition, we must have $n \leq m$. In our model, a rationality bound of n and cognitive bound of m are captured by a λ with $\lambda_k = 1$ for $k \leq n$, $\lambda_k \in (0, 1)$ for $n < k \leq m$, and $\lambda_k = 0$ for $k > m$.

A related distinction between rationality and cognitive ability was analyzed in [Alaoui and Penta \(2016\)](#). In that paper players choose whether to make the effort of reasoning as much as their cognitive bound allows. This idea is also similar to the framework in [Camerer, Loewenstein and Rabin \(2004\)](#), which unlike standard level- k reasoning, allows for uncertainty on the level of rationality attached to opponents. [Kets \(2014\)](#) and [Heifetz and Kets \(2016\)](#) generalize the σ -algebras attached to types so that they are able to capture a similar idea, and apply their construction to the study of the WY-discontinuity.

¹¹This has a similar flavor to “level- k reasoning,” with the distinction that level- k models begin with a level 0 type who takes a specific baseline action (possibly randomized), leading to specific actions for types at each level. We, rather, allow the full range of possible actions at stage 0 and continue with a set-valued concept at each stage. See [Stahl and Wilson \(1994\)](#) or [Nagel \(1995\)](#), among others, for “level- k reasoning”.

(v) *Unlimited depth of reasoning, with uncertainty on opponents' depth.* Pick sequence λ satisfying,

$$\forall n \in \mathbb{N}, \lambda_n \geq \lambda_{n+1}.$$

Here, we allow λ to be positive at all orders, which would signify that the player has unlimited depth of reasoning and attaches positive probability to all levels of opponents' reasoning. Again, λ_n is the probability he attaches to opponents' reasoning to at least depth n . He attaches probability $\lim_{k \rightarrow \infty} \lambda_k$ to his opponents' having unlimited depth of reasoning.

Most of the results of this paper (Propositions 1, 2 and 3, and Proposition 5) apply to every sequence λ , so in particular, also for the families of solution concepts considered above (in particular, Proposition 5 provides an epistemic foundation for all of them within a standard epistemic framework). Now, in Proposition 1 we will focus on a particular class of perturbations:

(vi) *Fading higher-order belief in rationality.*

$$\Lambda^0 = \left\{ \lambda \in \Lambda \left| \begin{array}{l} (i) \quad \lim_{n \rightarrow \infty} \lambda_n = 0, \\ (ii) \quad \lambda_n \geq \lambda_{n+1} \text{ for any } n \in \mathbb{N} \end{array} \right. \right\}.$$

The interpretation here is that each player is capable of reasoning to arbitrary levels, but is sufficiently uncertain of his opponents' depth that he loses almost all confidence at higher orders.

3.3 ELEMENTARY ROBUSTNESS PROPERTIES

Before continuing to our main results in Section 4, we present some elementary properties of ICR^λ . First, we check that ICR^λ is type-representation invariant; that is, the specific type structure employed to model a certain belief hierarchy does not affect interim correlated λ -rationalizable predictions. In order to perform a proper study of the problem let us introduce first a definition of ICR^λ in terms of types:

DEFINITION 2 (Interim correlated λ -rationalizability (Bayesian games)). *Let $\langle \mathcal{G}, \mathcal{T} \rangle$ be a Bayesian game and λ , a sequence of probabilities. Then, player i 's set of (interim correlated) λ -rationalizable actions for type t_i is defined as $\text{ICR}_i^{\lambda, \mathcal{T}}(t_i) = \bigcap_{n \geq 0} \text{ICR}_{i,n}^{\lambda, \mathcal{T}}(t_i)$, where $\text{ICR}_{i,0}^{\lambda, \mathcal{T}}(t_i) = A_i$ and $C_{i,0}^{\lambda, \mathcal{T}}(t_i) = \{\mu_i \in \Delta(T_{-i} \times A_{-i} \times \Theta) \mid \text{marg}_{T_{-i} \times \Theta} \mu_i = \pi_i(t_i)\}$, and recursively, for any $n \in \mathbb{N}$,*

$$\text{ICR}_{i,n}^{\lambda, \mathcal{T}}(t_i) = \left\{ a_i \in \text{ICR}_{i,n-1}^{\lambda, \mathcal{T}}(t_i) \mid a_i \in BR_i(\mu_i) \text{ for some } \mu_i \in C_{i,n-1}^{\lambda, \mathcal{T}}(t_i) \right\},$$

$$C_{i,n}^{\lambda, \mathcal{T}}(t_i) = \left\{ \mu_i \in C_{i,n-1}^{\lambda, \mathcal{T}}(t_i) \mid \mu_i[M] \geq \lambda_n \text{ for some measurable } M \subseteq \text{Graph} \left(\text{ICR}_{-i,n}^{\lambda, \mathcal{T}} \right) \times \Theta \right\}.$$

It is easy then to check that both definitions of ICR^λ are consistent: the set of ICR^λ actions corresponding to a type coincides with the set of ICR^λ corresponding to the belief hierarchy induced by the type:

PROPOSITION 1 (Type-representation invariance). *Let $\langle \mathcal{G}, \mathcal{T} \rangle$ be a Bayesian game. Then, for any player i , any type t_i and any sequence of probabilities λ , $\text{ICR}_i^{\lambda, \mathcal{T}}(t_i) = \text{ICR}_i^\lambda(\tau_i(t_i))$.*

Proposition 1 can be regarded as a robustness result of ICR^λ : different type representations of the same belief hierarchy lead to the same predictions. An additional robustness property of ICR^λ is presented in the following proposition, which shows that $\text{ICR}_i^\lambda : \mathcal{T}_i \rightrightarrows A_i$ is an upper-hemicontinuous correspondence. The latter implies that small misspecifications of belief hierarchies do not lead to originally unexpected behavior.

PROPOSITION 2 (Robustness to higher-order uncertainty about payoffs). *Let \mathcal{G} be a game with incomplete information. Then, for any $n \geq 0$, any player i , and any sequence of probabilities λ , correspondence $\text{ICR}_{i,n}^\lambda : \mathcal{T}_i \rightrightarrows A_i$ is upper-hemicontinuous. It follows that $\text{ICR}_i^\lambda : \mathcal{T}_i \rightrightarrows A_i$ is upper-hemicontinuous too.*

REMARK 1. *Notice that, for any $n \in \mathbb{N}$, any player i , and any sequence of probabilities λ , correspondence $\text{ICR}_{i,n}^\lambda : \mathcal{T}_i \rightrightarrows A_i$ has closed domain and is closed-valued; thus, it follows from Proposition 2 and the Closed Graph Theorem that $\text{Graph}(\text{ICR}_{-i,n}^\lambda)$ is closed and therefore, measurable. A useful consequence of this fact is that we can redefine ICR^λ as follows: player i 's set of λ -rationalizable actions for belief hierarchy τ_i is $\text{ICR}_i^\lambda(\tau_i) = \bigcap_{n \geq 0} \text{ICR}_{i,n}^\lambda(\tau_i)$, where $\text{ICR}_{i,0}^\lambda(\tau_i) = A_i$ and $C_{i,0}^\lambda(\tau_i) = \{\eta_i \in \Delta(\mathcal{T}_{-i} \times A_{-i} \times \Theta) \mid \text{marg}_{\mathcal{T}_{-i} \times \Theta} \eta_i = \varphi_i(\tau_i)\}$, and recursively, for any $n \in \mathbb{N}$,*

$$\begin{aligned} \text{ICR}_{i,n}^\lambda(\tau_i) &= \left\{ a_i \in \text{ICR}_{i,n-1}^\lambda(\tau_i) \mid a_i \in BR_i(\eta_i) \text{ for some } \eta_i \in C_{i,n-1}^\lambda(\tau_i) \right\}, \\ C_{i,n}^\lambda(\tau_i) &= \left\{ \eta_i \in C_{i,n-1}^\lambda(\tau_i) \mid \eta_i [\text{Graph}(\text{ICR}_{-i,n}^\lambda) \times \Theta] \geq \lambda_n \right\}. \end{aligned}$$

This identity in Remark 1 turns out to be helpful not only in simplifying the definition of ICR^λ , but also in the proof of our last result in this section, Proposition 3, which shows that: (i) ICR and ICR^λ coincide as perturbations in common

belief in rationality vanish (i.e., when $\lambda = \bar{1}$) and, based on the latter, that (ii) ICR is robust to higher-order uncertainty about rationality:¹² behavior is not only upper-hemicontinuous, but indeed, continuous (that is, also lower-hemicontinuous) when common belief in rationality is perturbed. Furthermore, Proposition 3, when combined with Propositions 1 and 2 above shows that both type-representation invariance and upper-hemicontinuity, as robustness properties of ICR, happen to be themselves robust to perturbations in common belief in rationality.

PROPOSITION 3 (Robustness to higher-order uncertainty about rationality). *Let \mathcal{G} be a game with incomplete information. Then, for any $n \in \mathbb{N}$, any player i and any belief hierarchy τ_i , it holds that $\text{ICR}_{i,n}(\tau_i) = \text{ICR}_{i,n}^{\bar{1}}(\tau_i)$. Furthermore, the correspondence given by $\lambda \mapsto \text{ICR}_{i,n}^\lambda(\tau_i)$ is upper-hemicontinuous everywhere and continuous at $\lambda = \bar{1}$. It follows that the correspondence given by $\lambda \mapsto \text{ICR}_i^\lambda(\tau_i)$ is upper-hemicontinuous everywhere and continuous at $\lambda = \bar{1}$ as well.*

4 MAIN RESULTS

We present now the main results of the paper, which study whether perturbations in higher-order belief in rationality eliminate the failures in continuity of rationalizability discovered by [Weinstein and Yildiz \(2007\)](#) in their Structure Theorem. To this end, we study the behavior of interim correlated λ -rationalizability for different λ . Our findings are twofold. Proposition 4 proves the robustness of the WY-discontinuity for constant sequences $\lambda = \bar{p}$; even under perturbation in common belief in rationality, if higher-order belief in rationality remains above some threshold, unique selection arguments à la [Weinstein and Yildiz \(2007\)](#) still work. However, Proposition 1 shows that the discontinuity goes away when λ converges to 0: if higher-order in rationality becomes eventually low enough, unique selection becomes impossible to accomplish. Similar results are found by [Heifetz and Kets \(2016\)](#), who instead of explicitly relaxing higher-order belief in rationality, introduce a more sophisticated framework that allows for higher-order uncertainty about players' cognitive bounds. The relation between [Heifetz and Kets's \(2016\)](#) work and this paper is examined in Section 4.2, where we also discuss the relevance of our results to global games.

¹²A related result by [Germano and Zuazo-Garin \(2017\)](#) shows that their notion of *p-rational outcomes* (which coincide with the correlated equilibria when $p = 1$ and otherwise generalize these by assuming common knowledge of mutual p -belief in rationality rather than common knowledge of rationality) are continuous in p , for any $p \leq 1$, which, in particular, implies robustness of correlated equilibria to bounded rationality.

In order to present Propositions 4 and 1, we need to recall first the definition of an auxiliary solution concept, interim strict correlated rationalizability (ICSR), which consists of a refinement of ICR where, at every elimination round, only *strict* best responses survive. More explicitly, given a game with incomplete information, \mathcal{G} , player i 's set of *interim correlated strictly rationalizable* actions for belief hierarchy τ_i is $\text{ICSR}_i(\tau_i) = \bigcap_{n \geq 0} \text{ICSR}_{i,n}(\tau_i)$, where $\text{ICSR}_{i,0}(\tau_i) = A_i$ and $D_{i,0}(\tau_i) = \{\eta_i \in \Delta(\mathcal{T}_{-i} \times A_{-i} \times \Theta) \mid \text{marg}_{\mathcal{T}_{-i} \times \Theta} \eta_i = \varphi_i(\tau_i)\}$, and then, recursively, for any $n \in \mathbb{N}$,

$$\text{ICSR}_{i,n}(\tau_i) = \{a_i \in A_i \mid BR_i(\eta_i) = \{a_i\} \text{ for some } \eta_i \in D_{i,n-1}(\tau_i)\},$$

$$D_{i,n}(\tau_i) = \{\eta_i \in D_{i,n-1}(\tau_i) \mid \eta_i[M] = 1 \text{ for some measurable } M \subseteq \text{Graph}(\text{ICSR}_{-i,n}) \times \Theta\}.$$

Again, a similar definition can be given when the original input is a Bayesian game $\langle \mathcal{G}, \mathcal{T} \rangle$ instead of a game with incomplete information: making the obvious changes player i 's set of strictly rationalizable actions for type t_i , $\text{ICSR}_i^{\mathcal{T}}(t_i)$, is obtained.

4.1 ROBUSTNESS OF THE WY-DISCONTINUITY

First, we show that the WY-discontinuity is robust to perturbations in common belief in rationality is perturbed that keep higher-order belief in rationality above some high enough threshold. This result will require [Weinstein and Yildiz's \(2007\)](#) richness condition:

DEFINITION 3 (Richness condition). *We say that a game with incomplete information satisfies the richness condition if for all actions a_i of any player i , there is a θ such that $u(a_i, a_{-i}, \theta) > u(a'_i, a_{-i}, \theta)$ for all (a'_i, a_{-i}) with $a'_i \neq a_i$.*

We can now state our first main result:

PROPOSITION 4 (Robustness of the WY-discontinuity). *Let $\langle \mathcal{G}, \mathcal{T} \rangle$ be a finite Bayesian game satisfying the richness condition. Then, for any player i , any type $t_i \in T_i$ and any $a_i \in \text{ICSR}_i^{\mathcal{T}}(t_i)$ there exists some $p < 1$ and some convergent sequence $(\tau_i^n)_{n \in \mathbb{N}}$ approaching t_i such that $\text{ICR}_i^{\bar{p}, \mathcal{T}}(\tau_i^n) = \{a_i\}$ for any $n \in \mathbb{N}$.*

Thus, for the sequence of types described in the WY-discontinuity, with unique rationalizable action, this property is robust to replacement of the ICR concept by $\text{ICR}^{\bar{p}}$. That is, even under this more permissive solution concept, representing bounded rationality, unique selection procedures work and any refinement sharper than ICSR will fail to be robust. In particular, any refinement which makes a selection among strict equilibria will fail to be robust. This is not the case if we allow

a different weakening of common belief in rationality, where belief in rationality becomes very low at high orders:

THEOREM 1 (Non-robustness of the WY-discontinuity). *Let \mathcal{G} be a game with incomplete information. Then for any belief hierarchy τ_i and any λ with $\lambda_n \rightarrow 0$ there exists a neighborhood U of τ_i such that $\text{ICSR}_i(\tau_i) \subseteq \text{ICR}_i^\lambda(\hat{\tau}_i)$ for any $\hat{\tau}_i \in U$.*

Proposition 1 implies that whenever the ICR and ICSR sets are identical at a hierarchy τ_i , ICR_i^λ will be continuous at τ_i for any λ with $\lambda_n \rightarrow 0$. Thus, while Proposition 4 states that the WY-discontinuity is robust to some kinds of perturbation in common belief in rationality, Proposition 1 tells us that it is not robust to every kind of perturbation: when the weight attached to higher-order belief in rationality becomes arbitrarily smaller as higher-order beliefs are considered, the unique selection of actions that can be made in the case of common belief in rationality turns out to be impossible. That is, as long as the assumption that higher-order beliefs become eventually negligible for players is introduced, no matter how slowly this diminishing impact of higher-order beliefs takes place, continuity of behavior with respect to perturbations of belief hierarchies is re-established (even when such perturbations are considered in the sense of the product topology).¹³ A special case was mentioned in Section 3.2: when $\lambda_n = 0$ for all $n > k$, a version of level- k reasoning. A rough justification for this result, in the level- k case, is that (1) ICSR actions remain in ICR_i^k when the first k levels of the hierarchy are close enough to the original type, and (2) the tail of the hierarchy becomes irrelevant when we reason only to level k .

A corollary to Proposition 1 states that the generic uniqueness result of [Weinstein and Yildiz](#) (closely related to the WY-discontinuity) also fails under the same conditions on λ .

COROLLARY 1 (Non-robustness of generic uniqueness). *Let \mathcal{G} be a game with incomplete information and finite set of action profiles. Then, for any player i for which there exists some τ_i such that $|\text{ICSR}_i(\tau_i)| > 1$, and for any λ with $\lambda_n \rightarrow 0$, the following set is not dense:*

$$\mathcal{U}_i^\lambda = \{ \tau_i \in \mathcal{T}_i \mid |\text{ICR}_i^\lambda(\tau_i)| = 1 \}.$$

¹³Notice that in one natural sense this is a small departure from common belief in rationality: if we put the product topology on the set of possible sequences λ , $\bar{1}$ is a limit point of the set of λ with $\lambda_n \rightarrow 0$ referenced in Proposition 1.

4.2 DISCUSSION

4.2.1 Implications for global games

Carlsson and van Damme (1993) introduced an argument for selection of “risk-dominant” equilibria, based on a discontinuity of the equilibrium correspondence. Given a complete-information game with multiple equilibria, they construct a family of incomplete-information games based on noisy observations of payoffs in the original game (a “global game”), where the risk-dominant equilibrium is unique even as the noise goes to zero. As discussed in Weinstein and Yildiz (2007), the WY-discontinuity weakens this argument in the sense that, for a larger family of perturbations of the original game, any action may be uniquely rationalizable. Our main results shed some light on these issues in the context of weakened common knowledge of rationality, as represented by ICR^λ . Under the conditions in Proposition 4, the selection of risk-dominant equilibria in global games will persist for large enough p , but so will the critique that the WY-discontinuity can lead to any selection. Under the conditions in Proposition 1, unique selection will be impossible for games close enough to the original game, because all ICSR actions (which include those played in a strict equilibrium, e.g. all actions in a 2×2 coordination game) remain rationalizable in small enough perturbations.

4.2.2 Almost common belief in rationality and almost common belief in infinite depth of reasoning

Proposition 1 shows that under certain arbitrarily *small* perturbations in common belief in rationality the WY-discontinuity vanishes; that is, continuity of behavior is restored under almost common belief in rationality. Going back to the terminology by Alaoui and Penta (2016) and Friedenberg, Kets and Kneeland (2016), the theorem departs from the standard model in Weinstein and Yildiz (2007) by introducing perturbations in common belief assumptions regarding players’ *rationality bounds*.

Strzalecki (2010) and Heifetz and Kets (2016) study the impact on the WY-discontinuity of perturbations in common belief assumptions regarding players’ *cognitive bound*. Specifically, Heifetz and Kets (2016) provide a framework that allows for modeling players’ uncertainty about each others’ depth of reasoning (i.e., cognitive bound), and show that under almost common belief in infinite depth of reasoning, the WY-discontinuity fails. That is, almost common belief in infinite depth is consistent with robust multiplicity (i.e., absence of generic uniqueness).

Insofar as the belief hierarchies employed in this paper are the same as the ones

in [Weinstein and Yildiz \(2007\)](#), they implicitly represent common belief in infinite depth of reasoning. Thus, the failure of the WY-discontinuity in Proposition 1, and, in particular, robust multiplicity, are consistent with common belief in infinite depth of reasoning. To see this, notice that Proposition 1 requires λ to be a sequence converging to 0, so that: (i) λ_n can be strictly greater than 0 for every n , implying that players' cognition is unbounded; however, (ii) there must exist some m such that $\lambda_n < 1$ for every $n \geq m$, meaning that every players' rationality bound is finite. Upon further scrutiny, such consistency is not surprising: while perturbing common belief in infinite depth of reasoning necessarily implies perturbing common belief in rationality, the opposite implication is not true. Since, given λ , the cognitive bound ($\sup\{n \in \mathbb{N} | \lambda_n > 0\}$) is by definition at least the rationality bound ($\sup\{n \in \mathbb{N} | \lambda_n = 1\}$), the existence of a cognitive bound necessarily implies the existence of a rationality bound. The opposite does not hold: a player with finite rationality bound might be cognitively unbounded, and therefore, failure of common belief in rationality does not imply failure of common belief in infinite depth or reasoning.

5 EPISTEMIC FOUNDATION OF λ -RATIONALIZABILITY

Finally, we formally analyze the epistemic foundation of interim correlated λ -rationalizability. The exercise corresponds to the incomplete information version of the case already studied by [Hu \(2007\)](#), with the addition that beliefs of different order can be given different consideration in the decision making process. Specifically, in Section 5.1 we introduce the epistemic framework needed for our study, which consists of a particular instance of the environment defined by [Battigalli, Di Tillio, Grillo and Penta \(2011\)](#). Next, in Section 5.2 we introduce the notion of common λ -belief, with λ a sequence of probabilities. This concept generalizes the standard notion of common p -belief due to [Monderer and Samet \(1987\)](#), allowing heterogeneous weights on higher-order beliefs. Common λ -belief serves as the base of our epistemic characterization result in Proposition 5, which generalizes several well-known characterization results in the Epistemic Game Theory literature.

5.1 EPISTEMIC FRAMEWORK

By applying [Brandenburger and Dekel's \(1993\)](#) construction to family of basic uncertainty spaces $(A_{-i} \times \Theta)_{i \in I}$, an alternative universal type space, $\langle \mathcal{E}_i, \psi_i \rangle_{i \in I}$, is obtained. We refer to each belief hierarchy $e_i \in \mathcal{E}_i$ as *epistemic hierarchy*. This way, following [Battigalli, Di Tillio, Grillo and Penta \(2011\)](#), 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 players' epistemic hierarchies and actions, and payoff states. The epistemic language is completed as follows:

5.1.1 Rationality and common (p -)belief

We say that player i is rational at state ω whenever her choice at ω is optimal given her first-order beliefs at ω . This event is formally represented by set $R_i = \{\omega \in \Omega \mid a_i(\omega) \in BR_i(e_{i,1}(\omega))\}$. As usual let $R = \bigcap_{i \in I} R_i$ and $R_{-i} = \bigcap_{i \in I} R_i$. Note that all these sets are closed and therefore measurable due to BR_i being closed-valued and Proj_{A_i} , continuous. Assumptions on players' beliefs can be represented by means of p -belief operators, as originally introduced by [Monderer and Samet \(1987\)](#). 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 ,

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

That is, event $B_i^p(E)$ is the collection of states in which player i assigns at least probability p to event E ; 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 it follows from the fact that every ψ_i is a homeomorphism that p -belief operators are closed-valued and therefore yield measurable sets. Finally, higher-order belief restrictions can be imposed using the *common p -belief operator*, which is recursively defined as follows: for each player i let $CB_i^p(E) = \bigcap_{n \geq 0} B_i^p(B^{n,p}(E))$, where $B^{0,p}(E) = E$, and recursively, $B^{n+1,p}(E) = B^p(B^{n,p}(E))$ for any $n \geq 0$. We write simply $CB_i(E) = CB_i^1(E)$ to represent *common belief*.

5.1.2 Epistemic hierarchies and belief hierarchies

Unsurprisingly, epistemic hierarchies and belief hierarchies are closely related. As shown by [Battigalli, Di Tillio, Grillo and Penta \(2011\)](#), it is possible to construct, by recursive marginalization, quotient maps $q_i : \mathcal{E}_i \rightarrow \mathcal{T}_i$ and $\bar{q}_i : \Delta(\mathcal{E}_{-i} \times A_{-i} \times \Theta) \rightarrow \Delta(\mathcal{T}_{-i} \times \Theta)$ that make the following diagram commutative:

$$\begin{array}{ccc}
\mathcal{E}_i & \overset{q_i}{\dashrightarrow} & \mathcal{T}_i \\
\psi_i \updownarrow & & \updownarrow \varphi_i \\
\Delta(\mathcal{E}_{-i} \times A_{-i} \times \Theta) & \overset{\bar{q}_i}{\dashrightarrow} & \Delta(\mathcal{T}_{-i} \times \Theta)
\end{array}$$

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 belief hierarchy τ_i , let $[q_i = \tau_i] = \{\omega \in \Omega \mid q_i(e_i(\omega)) = \tau_i\}$ be the event that player i 's belief hierarchy is exactly τ_i . Note that $[q_i = \tau_i]$ is closed due to q_i being continuous.

5.2 CHARACTERIZATION RESULT

We introduce now the epistemic operator that allows for our characterization result.

DEFINITION 4 (Common λ -belief). *Let $E \subseteq \Omega$ be an event, and λ , a sequence of probabilities. Let $B^{\lambda,0}(E) = E$, and set recursively $B^{\lambda,n+1}(E) = \bigcap_{i \in I} B_i^{\lambda,n+1}(B^{\lambda,n}(E))$ for each $n \geq 0$. Then, for each player i , $CB_i^\lambda(E) = \bigcap_{n \geq 0} B_i^{\lambda,n+1}(B^{\lambda,n}(E))$ is the event that player i exhibits common λ -belief in E .*

Thus, common λ -belief generalizes the notion of common p -belief so that at each iteration, the weight assigned to the corresponding epistemic restriction is not necessarily constant. The epistemic characterization of interim correlated λ -rationalizability exhibits then the expected pattern:

PROPOSITION 5 (Epistemic foundation of ICR^λ). *Let \mathcal{G} be a game with incomplete information, and λ , a sequence of probabilities. Then, interim correlated λ -rationalizability characterizes rationality and common λ -belief in rationality; i.e., for any player i and any belief hierarchy τ_i it holds that,*

$$ICR_i^\lambda(\tau_i) = \text{Proj}_{A_i} (R_i \cap CB_i^\lambda(R) \cap [q_i = \tau_i]).$$

The theoretical relevance of Proposition 5 lies in two features. First, as depicted in Figure 1, it shows that rationalizability is robust to a wide range of perturbations of common belief in rationality: not only perturbations à la p -belief, but also to the more general ones captured by non-constant λ parameters. This follows from the facts that: (i) interim correlated λ -rationalizability represents rational choice under departures from the standard rational benchmark by relaxing higher-order belief in rationality




Figure 1: Rationalizability and perturbations in common belief in rationality.

not necessarily weighting different order belief in an homogeneous way (Proposition 5) and (ii) interim correlated λ -rationalizability is upper-hemicontinuous on λ and indeed, continuous when $\lambda = \bar{\lambda}$ (Proposition 3). Second, since the result holds for arbitrary sequence λ , the epistemic foundation result covers the cases of particular λ sequences characterizing the different solution concepts reviewed in Section 2.3. This is already known in the case of standard solution concepts such as ICR (see Theorem 1 by Battigalli, Di Tillio, Grillo and Penta (2011), which corresponds to the $\lambda = \bar{\lambda}$ case) or p -rationalizability (see Proposition 1 by Hu (2007), which corresponds to the case of $\lambda = \bar{p}$ and τ_i exhibiting common belief in some game). The fact that solution concepts based on complex formal departures such as finite depth of reasoning models can be formalized and given epistemic formulation by means of already well-known tools reinforces the strength of the standard and classic game-theoretical approach.

A PROOFS: PROPERTIES OF ICR^λ

A.1 ELEMENTARY ROBUSTNESS PROPERTIES

For convenience, we begin with the proof of Proposition 2. An immediate corollary of this result is the characterization in Remark 1, which greatly simplifies the proof of Lemma 1, an auxiliary result presented below, and convenient itself for the proofs of Propositions 1 and 3.

PROPOSITION 2 (Robustness to higher-order uncertainty about payoffs). *Let \mathcal{G} be a game with incomplete information. Then, for any $n \geq 0$, any player i , and any sequence of probabilities λ , correspondence $\text{ICR}_{i,n}^\lambda : \mathcal{T}_i \rightrightarrows A_i$ is upper-hemicontinuous. It follows that $\text{ICR}_i^\lambda : \mathcal{T}_i \rightrightarrows A_i$ is upper-hemicontinuous too.*

Proof. We proceed by induction. The initial step (the $n = 0$) is immediate: $\tau_i \mapsto \text{ICR}_{i,0}^\lambda(\tau_i) = A_i$ is trivially upper-hemicontinuous for any $i \in I$ and $\lambda \in \Lambda$. For the inductive step, suppose that $n \geq 0$ is such that the claim holds. Then, to check the $(n + 1)^{\text{th}}$ case, fix $i \in I$ and $\lambda \in \Lambda$ and pick convergent sequence $(\tau_i^k)_{k \in \mathbb{N}}$ with limit τ_i and $a_i \in A_i$ such that $a_i \in \text{ICR}_{i,n+1}^\lambda(\tau_i^k)$ for any $k \in \mathbb{N}$. Then, we know that for any $k \in \mathbb{N}$ there is some $\eta_i^k \in C_{i,n}^\lambda(\tau_i^k)$ such that $a_i \in BR_i(\eta_i^k)$. Let $(\eta_i^{k_m})_{m \in \mathbb{N}}$ be a convergent subsequence of $(\eta_i^k)_{k \in \mathbb{N}}$ and let η_i denote its limit. Since $\text{marg}_{\mathcal{T}_{-i} \times \Theta}$ is continuous, $\eta_i \in C_{i,0}^\lambda(\tau_i)$. Now, notice that we know by the induction hypothesis that $\text{ICR}_{-i,\ell}^\lambda : \mathcal{T}_{-i} \rightrightarrows A_{-i}$ is upper-hemicontinuous for any $\ell = 1, \dots, n$. Then, it follows from the Closed Graph Theorem that for any $\ell = 1, \dots, n$, $M_\ell = \text{Graph}(\text{ICR}_{-i,\ell}^\lambda)$ is closed and therefore, measurable. Obviously, this implies that $\eta_i^{k_m}[M_\ell] \geq \lambda_\ell$ for any $\ell = 1, \dots, n$ and any $m \in \mathbb{N}$. Then, since $(\eta_i)_{m \in \mathbb{N}}^{k_m}$ converges to η_i and $(M_\ell)_{\ell=1}^n$ is a family of closed sets,

$$\eta_i[M_\ell] \geq \limsup_{m \rightarrow \infty} \eta_i^{k_m}[M_\ell] \geq \lambda_\ell$$

for any $\ell = 1, \dots, n$, and therefore, $\eta_i \in C_{i,n}^\lambda(\tau_i)$. Finally, the fact that BR_i is upper-hemicontinuous and $a_i \in BR_i(\eta_i^{k_m})$ for any $m \in \mathbb{N}$ implies that $a_i \in \text{ICR}_{i,n+1}^\lambda(\tau_i)$. ■

As said, the following characterization result, based on the fact that the graphs induced by interim correlated λ -rationalizability of finite order are measurable (consequence of Proposition 2), will be found useful in for the proofs of Propositions 1 and 3:

LEMMA 1. *Let \mathcal{G} be a game with incomplete information and λ , sequence of probabilities, and let $\lambda_0 = 1$. Then, for $n \in \mathbb{N}$, any player i and any belief hierarchy τ_i it*

holds that:

$$\text{ICR}_{i,n}^\lambda(\tau_i) =$$

$$= \left\{ a_i \in A_i \left| \begin{array}{l} \text{There exists some measurable } \sigma_{-i} : \mathcal{T}_{-i} \times \Theta \rightarrow \Delta(A_{-i}) \text{ such that:} \\ (i) \int_{\mathcal{T}_{-i} \times \Theta} \sigma_{-i}(\tau_{-i}, \theta) [\text{ICR}_{-i,k}^\lambda(\tau_{-i})] d\varphi_i(\tau_i) \geq \lambda_k \text{ for any } k = 1, \dots, n-1 \\ (ii) a_i \in \arg \max_{a'_i \in A_i} \int_{\mathcal{T}_{-i} \times \Theta} \left(\sum_{a_{-i} \in A_{-i}} \sigma_{-i}(\tau_{-i}, \theta)[a_{-i}] \cdot u_i((a_{-i}, a'_i), \theta) \right) d\varphi_i(\tau_i) \end{array} \right. \right\}.$$

Proof. We proceed by induction on n :

INITIAL STEP ($n = 1$). For the right-hand inclusion, pick $a_i \in \text{ICR}_{i,1}^\lambda(\tau_i)$ and $\eta_i \in \mathcal{C}_{i,0}^\lambda(\tau_i)$ such that $a_i \in BR_i(\eta_i)$. Since $\text{Proj}_{\mathcal{T}_{-i} \times \Theta} : \mathcal{T}_{-i} \times A_{-i} \times \Theta \rightarrow \mathcal{T}_{-i} \times \Theta$ is continuous and $\varphi_i(\tau_i)[E] = \eta_i[\text{Proj}_{\mathcal{T}_{-i} \times \Theta}^{-1}(E)]$ for any measurable $E \subseteq \mathcal{T}_{-i} \times \Theta$, then it follows immediately from the Disintegration Theorem that there exists a map $\sigma_{-i} : \mathcal{T}_{-i} \times \Theta \rightarrow \Delta(A_{-i})$ such that,¹⁴

(a) For each $E \subseteq A_{-i}$, map $\sigma_{-i}^E : \mathcal{T}_{-i} \times \Theta \rightarrow [0, 1]$ given $(\tau_{-i}, \theta) \mapsto \sigma_{-i}(\tau_{-i}, \theta)[E]$ is measurable. Hence, σ_{-i} is measurable too.¹⁵

(b) For any measurable $E \subseteq \mathcal{T}_{-i} \times A_{-i} \times \Theta$,

$$\mu_i[E] = \int_{\mathcal{T}_{-i} \times \Theta} \sigma_{-i}(\tau_{-i}, \theta)[\text{Proj}_{A_{-i}}(E \cap \{(\tau_{-i}, \theta)\} \times A_{-i})] d\varphi_i(\tau_i).$$

(c) We have that:

$$\begin{aligned} \int_{A_{-i} \times \Theta} u_i((a_{-i}; a_i), \theta) d(\text{marg}_{A_{-i} \times \Theta} \eta_i) &= \\ &= \int_{\mathcal{T}_{-i} \times \Theta} \left(\sum_{a_{-i} \in A_{-i}} \sigma_{-i}(\tau_{-i}, \theta)[a_{-i}] \cdot u_i((a_{-i}, a_i), \theta) \right) d\varphi_i(\tau_i). \end{aligned}$$

Then, since $\text{ICR}_{-i,0}^\lambda(\tau_{-i}) = A_{-i}$ for and $\tau_{-i} \in \mathcal{T}_{-i}$, σ_{-i} obviously satisfies conditions

¹⁴See Theorem 5.3.1 in [Ambrosio, Gigli and Savaré \(2006\)](#), p. 121. We are working with compact and metrizable spaces; thus, in particular, all of them are Polish and hence, Radon.

¹⁵Remember that we know from Lemma 4.5 by [Heifetz and Samet \(1998\)](#) that the Borel σ -algebra in corresponding to A_{-i} is generated by family $\{\{\mu_i \in \Delta(A_{-i}) | \mu_i[E] \geq p\} | E \subseteq A_{-i} \text{ and } p \in [0, 1]\}$. Hence, it follows from the measurability of each σ_{-i}^E that $\{(\tau_{-i}, \theta) \in \mathcal{T}_{-i} \times \Theta | \sigma_{-i}(\tau_{-i}, \theta)[E] \geq p\}$ is measurable for every $E \subseteq A_{-i}$ and every $p \in [0, 1]$. In consequence, σ_{-i} is measurable.

(i) and (ii) in the statement of the Lemma. For the left-hand inclusion, pick $a_i \in A_i$ and measurable $\sigma_{-i} : \mathcal{T}_{-i} \times \Theta \rightarrow \Delta(A_{-i})$ satisfying conditions (i) and (ii) above. Then, define measure $\eta_i \in \Delta(\mathcal{T}_{-i} \times A_{-i} \times \Theta)$ as follows:

$$\eta_i[E] = \int_{\mathcal{T}_{-i} \times \Theta} \left(\sum_{a_{-i} \in A_{-i}} \sigma_{-i}(\tau_{-i}, \theta) [\text{Proj}_{A_{-i}}(E \cap \{(\tau_{-i}, a_{-i}, \theta)\})] \right) d\varphi_i(\tau_i),$$

for any measurable $E \subseteq \mathcal{T}_{-i} \times A_{-i} \times \Theta$.¹⁶ We claim now that the following two hold:

- $\eta_i \in C_{i,0}^\lambda(\tau_i)$. To see it, pick measurable $E \subseteq \mathcal{T}_{-i} \times \Theta$ and develop:

$$\begin{aligned} \eta_i[E \times A_{-i}] &= \int_{\mathcal{T}_{-i} \times \Theta} \left(\sum_{a_{-i} \in A_{-i}} \sigma_{-i}(\tau_{-i}, \theta) [\text{Proj}_{A_{-i}}(E \times A_{-i} \cap \{(\tau_{-i}, a_{-i}, \theta)\})] \right) d\varphi_i(\tau_i) \\ &= \int_E \sigma_{-i}(\tau_{-i}, \theta)[A_{-i}] d\varphi_i(\tau_i) = \varphi_i(\tau_i)[E]. \end{aligned}$$

- $a_i \in BR_i(\eta_i)$. To see it, first, define, for each $a_{-i} \in A_{-i}$, measure $\nu_i(a_{-i}) \in \Delta(\mathcal{T}_{-i} \times \Theta)$ as $E \mapsto \eta_i[E \times \{a_{-i}\}]$. Then, for any $a'_i \in A_i$,

$$\begin{aligned} \int_{A_{-i} \times \Theta} u_i((a_{-i}; a'_i), \theta) d(\text{marg}_{A_{-i} \times \Theta} \eta_i) &= \\ &= \sum_{a_{-i} \in A_{-i}} \int_{\mathcal{T}_{-i} \times \Theta} u_i((a_{-i}; a'_i), \theta) d\nu_i(a_{-i}) \\ &= \int_{\mathcal{T}_{-i} \times \Theta} \left(\sum_{a_{-i} \in A_{-i}} \sigma_{-i}(\tau_{-i}, \theta)[a_{-i}] \cdot u_i((a_{-i}, a'_i), \theta) \right) d\varphi_i(\tau_i). \end{aligned}$$

Then, the fact that σ_{-i} satisfies property (ii) above proves the claim.

In consequence, $a_i \in \text{ICR}_{i,1}^\lambda(\tau_i)$.

INDUCTIVE STEP. Suppose that $n \geq 1$ is such that the claim holds. Let's check the $(n+1)$ th case. For the right-hand inclusion, pick $a_i \in \text{ICR}_{i,n+1}^\lambda(\tau_i)$ and $\eta_i \in C_{i,n}^\lambda(\tau_i)$ such that $a_i \in BR_i(\eta_i)$, and family $(M_k)_{k=1}^n$ of measurable sets such that $M_k \subseteq \text{Graph}(\text{ICR}_{-i,k}^\lambda)$ and $\eta_i[M_k] \geq \lambda_k$ for any $k = 1, \dots, n$. Then, since map $\text{Proj}_{\mathcal{T}_{-i} \times \Theta} : \mathcal{T}_{-i} \times A_{-i} \times \Theta \rightarrow \mathcal{T}_{-i} \times \Theta$ is continuous and $\varphi_i(\tau_i)[E] = \eta_i[\text{Proj}_{\mathcal{T}_{-i} \times \Theta}^{-1}(E)]$

¹⁶A similar argument to the one in the previous footnote proves that if σ_{-i} is measurable, then so is σ_{-i}^E for measurable set E . Since every set $\text{Proj}_{A_{-i}}(E \cap \{(\tau_{-i}, a_{-i}, \theta)\})$ is measurable, we conclude that η_i is a well-defined measure.

for any measurable $E \subseteq \mathcal{T}_{-i} \times \Theta$, we know again from the Disintegration Theorem that there exists a map $\sigma_{-i} : \mathcal{T}_{-i} \times \Theta \rightarrow \Delta(A_{-i})$ that satisfies properties (a), (b) and (c) in the paragraph above (in particular, we saw that such σ_{-i} is measurable). Condition (ii) in the statement of the lemma is trivially satisfied. To see (i), simply note that for any $k = 1, \dots, n$,

$$\begin{aligned} & \int_{\mathcal{T}_{-i} \times \Theta} \sigma_{-i}(\tau_{-i}, \theta) [\text{ICR}_{-i,k}^\lambda(\tau_{-i})] d\varphi_i(\tau_i) = \\ &= \int_{\mathcal{T}_{-i} \times \Theta} \sigma_{-i}(\tau_{-i}, \theta) [\text{Proj}_{A_{-i}}(\mathcal{T}_{-i} \times \text{ICR}_{-i,k}^\lambda(\tau_{-i}) \times \Theta \cap \{(\tau_{-i}, \theta)\} \times A_{-i})] d\varphi_i(\tau_i) \\ &\geq \int_{\mathcal{T}_{-i} \times \Theta} \sigma_{-i}(\tau_{-i}, \theta) [\text{Proj}_{A_{-i}}(M_k \cap \{(\tau_{-i}, \theta)\} \times A_{-i})] d\varphi_i(\tau_i) \\ &= \mu_i[M_k] \geq \lambda_k. \end{aligned}$$

For the left-hand inclusion, pick $a_i \in A_i$ and measurable map $\sigma_{-i} : \mathcal{T}_{-i} \times \Theta \rightarrow \Delta(A_{-i})$ satisfying conditions (i) and (ii) for the $(n+1)^{\text{th}}$ version of the statement of the lemma. Then, define measure $\eta_i \in \Delta(\mathcal{T}_{-i} \times A_{-i} \times \Theta)$ as follows:

$$\eta_i[E] = \int_{\mathcal{T}_{-i} \times \Theta} \left(\sum_{a_{-i} \in A_{-i}} \sigma_{-i}(\tau_{-i}, \theta) [\text{Proj}_{A_{-i}}(E \cap \{(\tau_{-i}, a_{-i}, \theta)\})] \right) d\varphi_i(\tau_i),$$

for any measurable $E \subseteq \mathcal{T}_{-i} \times A_{-i} \times \Theta$. We claim now that the following three hold:

- $\eta_i \in C_{i,0}^\lambda(\tau_i)$. To see it, pick measurable $E \subseteq \mathcal{T}_{-i} \times \Theta$ and develop:

$$\begin{aligned} \eta_i[E \times A_{-i}] &= \int_{\mathcal{T}_{-i} \times \Theta} \left(\sum_{a_{-i} \in A_{-i}} \sigma_{-i}(\tau_{-i}, \theta) [\text{Proj}_{A_{-i}}(E \times A_{-i} \cap \{(\tau_{-i}, a_{-i}, \theta)\})] \right) d\varphi_i(\tau_i) \\ &= \int_E \sigma_i(\tau_{-i}, \theta) [A_{-i}] d\varphi_i(\tau_i) = \varphi_i(\tau_i)[E]. \end{aligned}$$

- $\eta_i \in C_{i,n}^\lambda(\tau_i)$. Note that we know from Proposition 2 that $M_k = \text{Graph}(\text{ICR}_{-i,k}^\lambda)$ is measurable for any $k = 1, \dots, n$. Thus:

$$\begin{aligned} \eta_i[M_k] &= \int_{\mathcal{T}_{-i} \times \Theta} \sigma_{-i}(\tau_{-i}, \theta) [\text{Proj}_{A_{-i}}(M_k \cap \{(\tau_{-i}, \theta)\} \times A_{-i})] d\varphi_i(\tau_i) \\ &= \int_{\mathcal{T}_{-i} \times \Theta} \sigma_{-i}(\tau_{-i}, \theta) [\text{ICR}_{-i,k}^\lambda(\tau_{-i})] d\varphi_i(\tau_i) \geq \lambda_k \end{aligned}$$

for any $k = 1, \dots, n$.

- $a_i \in BR_i(\eta_i)$. To see it, first, define, for each $a_{-i} \in A_{-i}$, measure $\nu_i(a_{-i}) \in \Delta(\mathcal{T}_{-i} \times \Theta)$ given by $E \mapsto \eta_i[E \times \{a_{-i}\}]$. Then, for any $a'_i \in A_i$,

$$\begin{aligned} \int_{A_{-i} \times \Theta} u_i((a_{-i}; a'_i), \theta) d(\text{marg}_{A_{-i} \times \Theta} \eta_i) &= \\ &= \sum_{a_{-i} \in A_{-i}} \int_{\mathcal{T}_{-i} \times \Theta} u_i((a_{-i}; a'_i), \theta) d\nu_i(a_{-i}) \\ &= \int_{\mathcal{T}_{-i} \times \Theta} \left(\sum_{a_{-i} \in A_{-i}} \sigma_{-i}(\tau_{-i}, \theta)[a_{-i}] \cdot u_i((a_{-i}, a'_i), \theta) \right) d\varphi_i(\tau_i). \end{aligned}$$

The fact that σ_{-i} satisfies property (ii) above proves the claim.

This way, we conclude that $a_i \in \text{ICR}_{i,n+1}^\lambda(\tau_i)$. ■

We apply now Lemma 1 to the proofs of the two remaining propositions of Section 3.3:

PROPOSITION 1 (Type-representation invariance). *Let $\langle \mathcal{G}, \mathcal{F} \rangle$ be a Bayesian game. Then, for any player i , any type t_i and any sequence of probabilities λ , $\text{ICR}_i^{\lambda, \mathcal{F}}(t_i) = \text{ICR}_i^\lambda(\tau_i(t_i))$.*

Proof. We will prove the slightly more general claim: for any player i , any type t_i , any sequence of probabilities λ and any non-negative integer n , it holds that $\text{ICR}_{i,n}^{\lambda, \mathcal{F}}(t_i) = \text{ICR}_{i,n}^\lambda(\tau_i(t_i))$. Let's proceed by induction on n . The initial case ($n = 0$) holds trivially. For the inductive step, suppose that $n \geq 0$ is such that the claim holds for any $k = 0, \dots, n$, and fix $i \in I$, $t_i \in T_i$ and $\lambda \in \Lambda$. For the right inclusion, pick $a_i \in \text{ICR}_{i,n+1}^{\lambda, \mathcal{F}}(t_i)$ and $\mu_i \in C_{i,n}^{\lambda, \mathcal{F}}(t_i)$ such that $a_i \in BR_i(\mu_i)$, and for each $k = 1, \dots, n$, $M_k \subseteq \text{Graph}(\text{ICR}_{-i,k}^{\lambda, \mathcal{F}})$ such that $\mu_i[M_k] \geq \lambda_k$. Define now $\eta_i(\mu_i) \in \Delta(\mathcal{T}_{-i} \times A_{-i} \times \Theta)$ as follows:

$$E \mapsto \eta_i(\mu_i)[E] = \mu_i[\{(t_{-i}, a_{-i}, \theta) \in T_{-i} \times A_{-i} \times \Theta \mid ((\tau_{-i}(t_{-i})), a_{-i}, \theta) \in E\}],$$

for any measurable $E \subseteq \mathcal{T}_{-i} \times A_{-i} \times \Theta$. Since τ_{-i} is continuous, $\eta_i(\mu_i)$ is well-defined.¹⁷ Notice that we have (i) that $\text{marg}_{A_{-i} \times \Theta} \eta_i(\mu_i) = \text{marg}_{A_{-i} \times \Theta} \mu_i$ and (ii) that,¹⁸

$$\text{marg}_{\mathcal{T}_{-i} \times \Theta} \eta_i(\mu_i)[E] = \eta_i(\mu_i)[A_{-i} \times E]$$

¹⁷Due to every $\{(t_{-i}, a_{-i}, \theta) \in T_{-i} \times A_{-i} \times \Theta \mid ((\tau_{-i}(t_{-i})), a_{-i}, \theta) \in E\}$ being measurable.

¹⁸The fifth equality is a special case of formula (4) in Battigalli, Di Tillio, Grillo and Penta (2011), p.10.

$$\begin{aligned}
&= \mu_i[\{(t_{-i}, a_{-i}, \theta) \in T_{-i} \times A_{-i} \times \Theta \mid ((\tau_{-i}(t_{-i}), a_{-i}, \theta) \in A_{-i} \times E)\}] \\
&= \text{marg}_{T_{-i} \times \Theta} \mu_i[\{(t_{-i}, \theta) \in T_{-i} \times \Theta \mid (\tau_{-i}(t_{-i}), \theta) \in E\}] \\
&= \pi_i(t_i)[\{(t_{-i}, \theta) \in T_{-i} \times \Theta \mid (\tau_{-i}(t_{-i}), \theta) \in E\}] \\
&= \varphi_i(\tau_i(t_i))[E].
\end{aligned}$$

Thus, it follows from (i) that $a_i \in BR_i(\eta_i(\mu_i))$, and from (ii), that $\eta_i(\mu_i) \in C_{i,0}^\lambda(\tau_i(t_i))$. Now, fix $k = 0, \dots, n$ and note that we know, due to the induction hypothesis that,¹⁹

$$\begin{aligned}
\eta_i(\mu_i) [\text{Graph}(\text{ICR}_{-i,k}^\lambda) \times \Theta] &\geq \mu_i [\{(t_{-i}, a_{-i}) \in T_{-i} \times A_{-i} \mid a_{-i} \in \text{ICR}_{-i,k}^\lambda(\tau_{-i}(t_{-i}))\} \times \Theta] \\
&= \mu_i [\text{Graph}(\text{ICR}_{-i,k}^{\lambda, \mathcal{F}}) \times \Theta] \\
&\geq \mu_i [M_k] \geq \lambda_k.
\end{aligned}$$

Thus, we conclude that $\eta_i(\mu_i) \in C_{i,n}^\lambda(\tau_i(t_i))$. For the left inclusion we make use of Lemma 1. Pick $a_i \in \text{ICR}_{i,n+1}^\lambda(\tau_i(t_i))$ and measurable $\sigma_{-i} : \mathcal{T}_{-i} \times \Theta \rightarrow \Delta(A_{-i})$ satisfying conditions (i) and (ii) in the statement of the lemma. Since map $f_{-i} : T_{-i} \times \Theta \rightarrow \mathcal{T}_{-i} \times \Theta$ given by $(t_{-i}, \theta) \mapsto (\tau_{-i}(t_{-i}), \theta)$ is continuous, $\hat{\sigma}_{-i} = \sigma_{-i} \circ f_{-i}$ is measurable. We can then define $\mu_i \in \Delta(T_{-i} \times A_{-i} \times \Theta)$ as follows:

$$E \mapsto \mu_i[E] = \sum_{a_{-i} \in A_{-i}} \int_{T_{-i} \times \Theta} \hat{\sigma}_{-i}(t_{-i}, \theta) [E \cap \{(t_{-i}, a_{-i}, \theta)\}] d\pi_i(t_i),$$

for any measurable $E \subseteq T_{-i} \times A_{-i} \times \Theta$. Then, we have that:

- $\mu_i \in C_{i,0}^{\lambda, \mathcal{F}}(t_i)$. To see it, pick measurable $E \subseteq T_{-i} \times \Theta$ and develop:

$$\begin{aligned}
\mu_i[E \times A_{-i}] &= \int_{T_{-i} \times \Theta} \left(\sum_{a_{-i} \in A_{-i}} \hat{\sigma}_{-i}(t_{-i}, \theta) [\text{Proj}_{A_{-i}}(E \times A_{-i} \cap \{(t_{-i}, a_{-i}, \theta)\})] \right) d\pi_i(t_i) \\
&= \int_E \hat{\sigma}_{-i}(t_{-i}, \theta) [A_{-i}] d\pi_i(t_i) = \pi_i(t_i)[E].
\end{aligned}$$

- $\mu_i \in C_{i,n}^{\lambda, \mathcal{F}}(t_i)$. Consider continuous map $F_i : T_{-i} \times A_{-i} \times \Theta \mapsto \mathcal{T}_{-i} \times A_{-i} \times \Theta$ given by $(t_{-i}, a_{-i}, \theta) \mapsto (f_{-i}(t_{-i}, \theta), a_{-i})$. We know from the induction hypothesis that $M_k = \text{Graph}(\text{ICR}_{-i,k}^{\lambda, \mathcal{F}}) = F_i^{-1}(\text{Graph}(\text{ICR}_{-i,k}^\lambda))$ for any $k = 1, \dots, n$. It follows from Proposition 2 and the continuity of F_i that $(M_k)_{k=1}^n$ is a family of measurable sets which obviously satisfies that $M_k \subseteq \text{Graph}(\text{ICR}_{i,k}^{\lambda, \mathcal{F}})$ for any

¹⁹Each $\text{Graph}(\text{ICR}_{-i,k}^{\lambda, \mathcal{F}})$ is clearly measurable, see Footnote 17.

$k = 1, \dots, n$. Then, we have that:

$$\begin{aligned}\mu_i[M_k] &= \int_{T_{-i} \times \Theta} \sigma_{-i}(t_{-i}, \theta) [\text{Proj}_{A_{-i}}(M_k \cap \{(t_{-i}, \theta)\} \times A_{-i})] d\pi_i(t_i) \\ &= \int_{T_{-i} \times \Theta} \sigma_{-i}(t_{-i}, \theta) [\text{ICR}_{-i,k}^{\lambda, \mathcal{F}}(t_{-i})] d\pi_i(t_i) \geq \lambda_k\end{aligned}$$

for any $k = 1, \dots, n$.

- $a_i \in BR_i(\mu_i)$. Note first that:

$$\begin{aligned}& \int_{T_{-i} \times \Theta} \left(\sum_{a_{-i} \in A_{-i}} \hat{\sigma}_{-i}(t_{-i}, \theta) [a_{-i}] \cdot u_i((a_{-i}, a'_i), \theta) \right) d\pi_i(t_i) = \\ &= \int_{T_{-i} \times \Theta} \left(\int_{\tau_{-i}^{-1}(\tau_{-i}) \times \{\theta\}} \left(\sum_{a_{-i} \in A_{-i}} \hat{\sigma}_{-i}(t_{-i}, \theta) [a_{-i}] \cdot u_i((a_{-i}, a'_i), \theta) \right) d\pi_i(t_i) \right) d\varphi_i(\tau_i) \\ &= \int_{T_{-i} \times \Theta} \left(\int_{\tau_{-i}^{-1}(\tau_{-i}) \times \{\theta\}} \left(\sum_{a_{-i} \in A_{-i}} \sigma_{-i}(f_{-i}(t_{-i}), \theta) [a_{-i}] \cdot u_i((a_{-i}, a'_i), \theta) \right) d\pi_i(t_i) \right) d\varphi_i(\tau_i) \\ &= \int_{T_{-i} \times \Theta} \left(\int_{\tau_{-i}^{-1}(\tau_{-i}) \times \{\theta\}} \left(\sum_{a_{-i} \in A_{-i}} \sigma_{-i}(\tau_{-i}, \theta) [a_{-i}] \cdot u_i((a_{-i}, a'_i), \theta) \right) d\pi_i(t_i) \right) d\varphi_i(\tau_i) \\ &= \int_{T_{-i} \times \Theta} \left(\sum_{a_{-i} \in A_{-i}} \sigma_{-i}(\tau_{-i}, \theta) [a_{-i}] \cdot u_i((a_{-i}, a'_i), \theta) \right) d\varphi_i(\tau_i).\end{aligned}$$

Now, define for each $a_{-i} \in A_{-i}$ measure $\nu_i(a_{-i}) \in \Delta(T_{-i} \times \Theta)$ as $E \mapsto \mu_i[E \times \{a_{-i}\}]$. Then, for any $a'_i \in A_i$,

$$\begin{aligned}\int_{A_{-i} \times \Theta} u_i((a_{-i}; a'_i), \theta) d(\text{marg}_{A_{-i} \times \Theta} \mu_i) &= \\ &= \sum_{a_{-i} \in A_{-i}} \int_{T_{-i} \times \Theta} u_i((a_{-i}; a'_i), \theta) d\nu_i(a_{-i}) \\ &= \int_{T_{-i} \times \Theta} \left(\sum_{a_{-i} \in A_{-i}} \hat{\sigma}_{-i}(t_{-i}, \theta) [a_{-i}] \cdot u_i((a_{-i}, a'_i), \theta) \right) d\pi_i(t_i).\end{aligned}$$

Then, the fact that σ_{-i} satisfies property (ii) proves the claim.

Thus, we conclude that $a_i \in \text{ICR}_{i,n+1}^{\lambda, \mathcal{F}}(t_i)$. ■

PROPOSITION 3 (Robustness to higher-order uncertainty about rationality). *Let \mathcal{G} be a game with incomplete information. Then, for any $n \in \mathbb{N}$, any player i and any belief hierarchy τ_i , it holds that $\text{ICR}_{i,n}(\tau_i) = \text{ICR}_{i,n}^{\bar{1}}(\tau_i)$. Furthermore, the correspondence given by $\lambda \mapsto \text{ICR}_{i,n}^\lambda(\tau_i)$ is upper-hemicontinuous everywhere and continuous at $\lambda = \bar{1}$. It follows that the correspondence given by $\lambda \mapsto \text{ICR}_i^\lambda(\tau_i)$ is upper-hemicontinuous everywhere and continuous at $\lambda = \bar{1}$ as well.*

Proof. Since it follows immediately from Lemma 1 that for any $n \in \mathbb{N}$, any $i \in I$ and any $\tau_i \in \mathcal{T}_i$, $\text{ICR}_{i,n}(\tau_i) = \text{ICR}_{i,n}^1(\tau_i)$, we focus on the claims concerning continuity. We prove them separately:

UPPER-HEMICONTINUITY. We prove first the following claim: for any $i \in I$, any $\tau_i \in \mathcal{T}_i$ and any $n \geq 0$, correspondence $\lambda \mapsto \text{ICR}_{i,n}^\lambda(\tau_i)$ is upper-hemicontinuous. We proceed by induction on n . The initial step ($n = 0$) holds trivially. For the inductive step, suppose that $n \geq 0$ is such that the claim holds for any $k = 0, \dots, n$. In particular, note that each $\text{ICR}_{-i,k}^{(\cdot)}(\tau_{-i})$ is compact-valued, and hence, upper-hemicontinuity implies that $\bigcap_{n \geq 0} \text{ICR}_{-i,k}^{\lambda^n}(\tau_{-i}) \subseteq \text{ICR}_{-i,k}^\lambda(\tau_{-i})$ for any $(\lambda^n)_{n \in \mathbb{N}} \rightarrow \lambda$.²⁰ Now, fix $i \in I$ and $\tau_i \in \mathcal{T}_i$, pick convergent sequence $(\lambda^m, a_i^m)_{m \in \mathbb{N}} \subseteq \Lambda \times A_i$ such that $a_i^m \in \text{ICR}_{i,n+1}^{\lambda^m}(\tau_i)$ for any $m \in \mathbb{N}$, and denote by (λ, a_i) the limit of the sequence. We need to check that $a_i \in \text{ICR}_{i,n+1}^\lambda(\tau_i)$. First, pick $(\eta_i^m)_{m \in \mathbb{N}} \in \prod_{m \in \mathbb{N}} C_{i,n}^{\lambda^m}(\tau_i)$ such that $a_i^m \in BR_i(\eta_i^m)$ for any $m \in \mathbb{N}$, and notice that, since $\Delta(\mathcal{T}_{-i} \times A_{-i} \times \Theta)$ is compact, there exists a convergent subsequence $(\eta_i^{m_\ell})_{\ell \in \mathbb{N}} \subseteq (\eta_i^m)_{m \in \mathbb{N}}$ with limit η_i . Obviously, $(a_i^{m_\ell})_{\ell \in \mathbb{N}}$ converges to a_i , and thus, we know from the upper-hemicontinuity of BR_i that $a_i \in BR_i(\eta_i)$. Since $\text{marg}_{\mathcal{T}_{-i} \times \Theta}$ is continuous we also know that $\text{marg}_{\mathcal{T}_{-i} \times \Theta} \eta_i = \varphi_i(\tau_i)$.

It only remains to be checked that $\eta_i[\text{Graph}(\text{ICR}_{-i,k}^\lambda) \times \Theta] \geq \lambda_k$ for any $k = 0, \dots, n$. Fix $k = 0, \dots, n$ and notice that $\eta_i^{m_\ell}[\text{Graph}(\text{ICR}_{-i,k}^{\lambda^{m_\ell}}) \times \Theta] \geq \lambda_k^{m_\ell}$ for any $\ell \in \mathbb{N}$. Then, set $(\hat{\lambda}_k^{m_\ell}) = (\inf_{s \geq \ell} \lambda_k^{m_s})_{k \in \mathbb{N}}$ and $A_{k,m_\ell} = \bigcup_{t \geq \ell} \text{Graph}(\text{ICR}_{-i,k}^{\lambda^{m_t}})$ for any $\ell \in \mathbb{N}$. Since $(\hat{\lambda}_k^{m_\ell})$ is a weakly increasing sequence (i.e, for any $t \geq \ell$, $\hat{\lambda}_k^{m_t} \geq \hat{\lambda}_k^{m_\ell}$) and, clearly, $\lambda_k^{m_\ell} \geq \hat{\lambda}_k^{m_\ell}$, the following hold for any $\ell \in \mathbb{N}$,

- (i) $\text{Graph}(\text{ICR}_{-i,k}^{\hat{\lambda}_k^{m_\ell}}) \subseteq \text{Graph}(\text{ICR}_{-i,k}^{\lambda^{m_\ell}})$ for any $t \geq \ell$.
- (ii) $\text{Graph}(\text{ICR}_{-i,k}^{\lambda^{m_\ell}}) \subseteq \text{Graph}(\text{ICR}_{-i,k}^{\hat{\lambda}_k^{m_\ell}})$.

²⁰Just write: $\Gamma(\lambda) = \text{ICR}_{-i,k}^\lambda(\tau_{-i})$. Since Γ is compact-valued and upper-hemicontinuous, then $a_{-i} \in \Gamma(\lambda)$ for any $(\lambda^n)_{n \in \mathbb{N}}$ converging to λ such that $a_{-i} \in \Gamma(\lambda^n)$ for any $n \in \mathbb{N}$. Thus, $\bigcap_{n \in \mathbb{N}} \Gamma(\lambda^n) \subseteq \Gamma(\lambda)$.

It follows from (i) and (ii) that $\overline{A_{k,m_\ell}} \subseteq \text{Graph}(\text{ICR}_{-i,k}^{\hat{\lambda}^{m_\ell}})$ for any $\ell \in \mathbb{N}$.²¹ Now, notice that for any $\ell \in \mathbb{N}$, $\eta_i^{m_\ell+r}[\overline{A_{k,m_\ell}} \times \Theta] \geq \hat{\lambda}_k^{m_\ell}$, and that $(\eta_i^{m_\ell+r})_{r \geq 0}$ converges to η_i . Thus, we know from Theorem 15.3 by Aliprantis and Border (1999) that $\eta_i[\overline{A_{k,m_\ell}} \times \Theta] \geq \hat{\lambda}_k^{m_\ell}$, and therefore, that $\eta_i[\text{Graph}(\text{ICR}_{-i,k}^{\hat{\lambda}^{m_\ell}}) \times \Theta] \geq \hat{\lambda}_k^{m_\ell}$ for any $\ell \in \mathbb{N}$. The latter, together with (i) above and the fact that $(\hat{\lambda}_k^{m_\ell})_{\ell \in \mathbb{N}}$ converges to λ_k implies that,

$$\begin{aligned} \eta_i\left[\bigcap_{\ell \in \mathbb{N}} \text{Graph}(\text{ICR}_{-i,k}^{\hat{\lambda}^{m_\ell}}) \times \Theta\right] &= \eta_i\left[\lim_{\ell \rightarrow \infty} (\text{Graph}(\text{ICR}_{-i,k}^{\hat{\lambda}^{m_\ell}}) \times \Theta)\right] \\ &= \lim_{\ell \rightarrow \infty} \eta_i[\text{Graph}(\text{ICR}_{-i,k}^{\hat{\lambda}^{m_\ell}}) \times \Theta] \\ &\geq \lim_{\ell \rightarrow \infty} \hat{\lambda}_k^{m_\ell} = \lambda_k. \end{aligned}$$

Notice that we know from the induction hypothesis (see Footnote 17), again together with the fact that $(\hat{\lambda}_k^{m_\ell})_{\ell \in \mathbb{N}}$ converges to λ_k , that,

$$\eta_i[\text{Graph}(\text{ICR}_{-i,k}^\lambda) \times \Theta] \geq \eta_i\left[\bigcap_{\ell \in \mathbb{N}} \text{Graph}(\text{ICR}_{-i,k}^{\hat{\lambda}^{m_\ell}}) \times \Theta\right].$$

Thus, we conclude from the last two that $\eta_i[\text{Graph}(\text{ICR}_{-i,k}^\lambda) \times \Theta] \geq \lambda_k$ and hence, that $\eta_i \in C_{i,k}^\lambda(\tau_i)$, and $a_i \in \text{ICR}_{i,n+1}^\lambda(\tau_i)$.

It follows from the above that for any $i \in I$, any $\tau_i \in \mathcal{T}_i$ and any $n \geq 0$, $\text{Graph}(\text{ICR}_{i,n}^{(\cdot)}(\tau_i))$ is closed, and thus, that so is $\text{ICR}_{i,n}^\lambda(\tau_i) = \text{Proj}_{A_i}((\{\lambda\} \times A_i) \cap \text{Graph}(\text{ICR}_{i,n}^{(\cdot)}(\tau_i)))$ for any $\lambda \in \Lambda$. Thus, $\text{ICR}_{i,n}^{(\cdot)}(\tau_i)$ is a compact-valued correspondence, and hence, by Theorem 17.25 in Aliprantis and Border (1999), we conclude that $\text{ICR}_i^{(\cdot)}(\tau_i) = \bigcap_{n \geq 0} \text{ICR}_{i,n}^{(\cdot)}(\tau_i)$ is upper-hemicontinuous.

CONTINUITY AT $\lambda = 1$. Fix $i \in I$, and $\tau_i \in \mathcal{T}_i$. It suffices to check lower-hemicontinuity at $\lambda = 1$; 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 $\text{ICR}_i^\lambda(\tau_i) \cap U \neq \emptyset$, there exists a neighborhood $V \subseteq \Lambda$ of $\lambda = 1$ such that if $\lambda' \in V$, then $\text{ICR}_i^{\lambda'}(\tau_i) \cap U \neq \emptyset$. This follows immediately from the fact that, since $\lambda'_n \geq \lambda_n$ for any $n \in \mathbb{N}$, then $\text{ICR}_i^\lambda(\tau_i) \subseteq \text{ICR}_i^{\lambda'}(\tau_i)$. \blacksquare

²¹By $\overline{A_{k,m_\ell}}$ we denote the closure of A_{k,m_ℓ} ; note that we know from Proposition 2 that $\text{ICR}_{-i,k}^{\lambda_{m_\ell}}$ has closed graph.

A.2 EPISTEMIC CHARACTERIZATION

PROPOSITION 5 (Epistemic foundation of ICR^λ). *Let \mathcal{G} be a game with incomplete information, and λ , a sequence of probabilities. Then, interim correlated λ -rationalizability characterizes rationality and common λ -belief in rationality; i.e., for any player i and any belief hierarchy τ_i it holds that,*

$$\text{ICR}_i^\lambda(\tau_i) = \text{Proj}_{A_i} (R_i \cap CB_i^\lambda(R) \cap [q_i = \tau_i]).$$

Proof. Fix sequence of probabilities λ . Now, first, for any $i \in I$ and any $n \geq 1$ define auxiliary correspondence $\Phi_{i,n} : \text{Graph}(\text{ICR}_{i,n}^\lambda) \rightrightarrows \Omega_i$ as follows:

$$(\tau_i, a_i) \mapsto \{e_i \in q_i^{-1}(\tau_i) \mid (e_i, a_i) \in \text{Proj}_{\mathcal{E}_i \times A_i}(R_i \cap B_i^{\lambda, n-1}(R))\} \times \{a_i\}.$$

Note that for any $i \in I$ and $n \in \mathbb{N}$, correspondence $\Phi_{i,n-1}$ has closed graph: pick convergent sequence $(\tau_i^m, a_i^m, e_i^m, a_i^m)_{m \in \mathbb{N}} \subseteq \text{Graph}(\Phi_{i,n-1})$ with limit (τ_i, a_i, e_i, a_i) . Since $q_i(e_i^m) = \tau_i^m$ for any $m \in \mathbb{N}$ and q_i is continuous, we know that $e_i \in q_i^{-1}(\tau_i)$. Thus, it suffices to check that $(e_i, a_i) \in \text{Proj}_{\mathcal{E}_i \times A_i}(R_i \cap B_i^{\lambda, n-1}(R))$. But the latter is obvious: it follows immediately from the facts that $R_i \cap B_i^{\lambda, n-1}(R)$ is closed and $(e_i^m, a_i^m)_{m \in \mathbb{N}} \subseteq \text{Proj}_{\mathcal{E}_i \times A_i}(R_i \cap B_i^{\lambda, n-1}(R))$. This way, we conclude that $(\tau_i, a_i, e_i, a_i) \in \text{Graph}(\Phi_{i,n-1})$.

Now, for any $i \in I$ denote $B_i^{\lambda, 0}(R) = \Omega$. Let's prove that for any $n \geq 0$ we have that,

$$\text{Graph}(\text{ICR}_{i,n+1}^\lambda) = \{(q_i(e_i), a_i) \in \mathcal{T}_i \times A_i \mid (e_i, a_i) \in \text{Proj}_{\mathcal{E}_i \times A_i}(R_i \cap B_i^{\lambda, n}(R))\}.$$

We proceed by induction:

INITIAL STEP. Fix $i \in I$. For the left inclusion, pick $\omega \in R_i$ and set $(\tau_i, a_i) = (q_i(e_i(\omega)), a_i(\omega))$. Define now $\eta_i \in \Delta(\mathcal{T}_{-i} \times A_{-i} \times \Theta)$ as follows:

$$E \mapsto \eta_i[E] = \psi_i(e_i(\omega))[\{(e_{-i}, a_{-i}, \theta) \in \mathcal{E}_{-i} \times A_{-i} \times \Theta \mid (q_{-i}(e_{-i}), a_{-i}, \theta) \in E\}].$$

Since q_{-i} is a homeomorphism, η_i is well-defined, and obviously, it satisfies the following two conditions: (i) $\text{marg}_{\mathcal{T}_{-i} \times \Theta} \eta_i = q_i(\tau_i)$ and (ii) $\text{marg}_{A_{-i} \times \Theta} \eta_i = e_{i,1}(\omega)$. Thus, we have, first, that $\eta_i \in C_{i,0}^\lambda(\tau_i)$, and, second, since $\omega \in R_i$, that $a_i \in BR_i(\eta_i)$. In consequence, $(\tau_i, a_i) \in \text{Graph}(\text{ICR}_{i,1}^\lambda)$. For the right inclusion, define first correspondence $\Phi_{i,0} : \mathcal{T}_i \times A_i \rightrightarrows \Omega_i$ as follows: $(\tau_i, a_i) \mapsto q_i^{-1}(\tau_i) \times \{a_i\}$. Obviously, $\Phi_{i,0}$ is non-empty and has closed graph. Thus, it is also weakly measurable and then,

we know from the Kuratowski-Ryll Nardzewski Selection Theorem that it admits a measurable selector $\phi_{i,0}$. Let $\phi_{-i,0} = (\phi_{j,0})_{j \neq i}$. Next, pick $(\tau_i, a_i) \in \mathcal{T}_i \times A_i$ such that $a_i \in \text{ICR}_{i,1}^\lambda(\tau_i)$, and $\eta_i \in C_{i,0}^\lambda(\tau_i)$ such that $a_i \in BR_i(\eta_i)$, and define belief $\psi_i(\eta_i) \in \Delta(\mathcal{E}_{-i} \times A_{-i} \times \Theta)$ as follows:

$$E \mapsto \psi_i(\eta_i)[E] = \eta_i[\{(\tau_{-i}, a_{-i}, \theta) \in \mathcal{T}_{-i} \times A_{-i} \times \Theta \mid (\phi_{-i,0}(\tau_{-i}, a_{-i}), \theta) \in E\}].$$

Since $\phi_{-i,0}$ is measurable and its domain is $\mathcal{T}_{-i} \times A_{-i}$, $\psi_i(\eta_i)$ is well-defined. Set $e_i = \psi_i^{-1}(\psi_i(\eta_i))$. Then, we have that: (i) $\text{marg}_{\mathcal{T}_{-i} \times \Theta} \psi_i(e_i) = \text{marg}_{\mathcal{T}_{-i} \times \Theta} \eta_i$ and (ii) $e_{i,1} = \text{marg}_{A_{-i} \times \Theta} \eta_i$. Thus, it follows that $q_i(e_i) = \tau_i$ and $a_i \in BR_i(e_{i,1})$, and therefore, that $(e_i, a_i) \in \text{Proj}_{\mathcal{E}_i \times A_i}(R_i)$. Notice that, in particular, the proof of the right inclusion implies that $\Phi_{i,1}$ is non-empty.

INDUCTIVE STEP. Suppose that $n \geq 0$ is such that for any $k = 0, \dots, n$ the claim holds and $\Phi_{i,k+1}$ is non-empty for any $i \in I$. Fix $i \in I$. For the left inclusion, pick $\omega \in R_i \cap B_i^{\lambda, n+1}(R)$ and let $(\tau_i, a_i) = (q_i(e_i(\omega)), a_i(\omega))$. Define belief $\eta_i \in \Delta(\mathcal{T}_{-i} \times A_{-i} \times \Theta)$ as follows:

$$E \mapsto \eta_i[E] = \psi_i(e_i(\omega))[\{(e_{-i}, a_{-i}, \theta) \in \mathcal{E}_{-i} \times A_{-i} \times \Theta \mid (q_{-i}(e_{-i}), a_{-i}, \theta) \in E\}].$$

Since q_{-i} is a homeomorphism, η_i is well-defined, and clearly, it satisfies the following two conditions: (i) $\text{marg}_{\mathcal{T}_{-i} \times \Theta} \eta_i = q_i(\tau_i)$ and (ii) $\text{marg}_{A_{-i} \times \Theta} \eta_i = e_{i,1}(\omega)$. Thus, obviously, we have, first, that $\eta_i \in C_{i,0}^\lambda(\tau_i)$, and second, since $\omega \in R_i$, that $a_i \in BR_i(\eta_i)$. Finally, notice that, since $\omega \in B_i^{\lambda, n+1}(R)$, for any $k = 1, \dots, n+1$ it holds that,

$$\begin{aligned} & \eta_i[\text{Graph}(\text{ICR}_{-i,k}^\lambda) \times \Theta] = \\ & = \eta_i[\{q_{-i}(e_{-i}), a_{-i} \in \mathcal{T}_{-i} \times A_{-i} \mid (e_{-i}, a_{-i}) \in \text{Proj}_{\mathcal{E}_{-i} \times A_{-i}}(R_{-i} \cap B_{-i}^{\lambda, k-1}(R))\} \times \Theta] \\ & = \psi_i(e_i(\omega))[\{(e_{-i}, a_{-i}) \in \mathcal{E}_{-i} \times A_{-i} \mid a_{-i} \in \text{ICR}_{-i,k}^\lambda(q_{-i}(e_{-i}))\} \times \Theta] \\ & = \psi_i(e_i(\omega)) \left[\left\{ (e_{-i}, a_{-i}) \in \Omega_{-i} \mid \begin{array}{l} \text{There exists some } e'_i \in q_{-i}^{-1}(e_{-i}) \text{ such that} \\ (e'_{-i}, a_{-i}) \in \text{Proj}_{\mathcal{E}_{-i} \times A_{-i}}(R_{-i} \cap B_{-i}^{\lambda, k-1}(R)) \end{array} \right\} \times \Theta \right] \\ & \geq \psi_i(e_i(\omega))[\{(e_{-i}, a_{-i}, \theta) \in \Omega_{-i} \times \Theta \mid (e_{-i}, a_{-i}) \in \text{Proj}_{\mathcal{E}_{-i} \times A_{-i}}(R_{-i} \cap B_{-i}^{\lambda, k-1}(R))\}] \\ & = \psi_i(e_i(\omega))[\{(\omega'_{-i}, \theta) \in \Omega_{-i} \times \Theta \mid (\omega'_{-i}, \omega_i, \theta) \in R_{-i} \cap B_{-i}^{\lambda, k-1}(R)\}] \geq \lambda_k. \end{aligned}$$

Thus, $\eta_i \in C_{i,k}^\lambda(\tau_i)$ for any $k = 0, \dots, n+1$ and, in consequence, $(\tau_i, a_i) \in \text{Graph}(\text{ICR}_{i, n+2}^\lambda)$. For the right inclusion, pick $(\tau_i, a_i) \in \mathcal{T}_i \times A_i$ such that $a_i \in \text{ICR}_{i, n+2}^\lambda(\tau_i)$, and $\eta_i \in C_{i, n+1}^\lambda(\tau_i)$ such that $a_i \in BR_i(\eta_i)$. We know from the induction hypothesis that

$\Phi_{j,n+1}$ is non-empty for any $j \neq i$. Thus, since every $\Phi_{j,n+1}$ has closed graph, and hence, is weakly measurable, there exists a measurable map $\phi_{-i,n+1} = (\phi_{j,n+1})_{j \neq i}$ where for each $j \neq i$ map $\phi_{j,n+1}$ is a measurable selector of $\Phi_{j,n+1}$. Next, let's introduce the following notational convention: let $Z_{-i,k} = \text{Graph}(\text{ICR}_{-i,k}^\lambda)$ and $W_{-i,k} = \text{Proj}_{\Omega}(R_{-i} \cap B_{-i}^{\lambda,k}(R))$ for any $k = 0, \dots, n+1$. Then, define $\psi_i(\eta_i) \in \Delta(\mathcal{E}_{-i} \times A_{-i} \times \Theta)$ as follows:

$$E \mapsto \psi_i(\eta_i)[E] = \sum_{k=0}^{n+1} \psi_i^k(\eta_i)[E],$$

where,

$$\begin{aligned} \psi_i^{n+1}(\eta_i)[E] &= \eta_i[\{(\tau_{-i}, a_{-i}, \theta) \in Z_{-i,n+1} \times \Theta \mid (\phi_{-i,n+1}(\tau_{-i}, a_{-i}), \theta) \in E\}], \text{ and,} \\ \psi_i^k(\eta_i)[E] &= \eta_i[\{(\tau_{-i}, a_{-i}, \theta) \in (Z_{-i,k} \setminus Z_{-i,k+1}) \times \Theta \mid (\phi_{-i,k}(\tau_{-i}, a_{-i}), \theta) \in E\}], \end{aligned}$$

for any $k = 0, \dots, n$. Since every $\phi_{-i,k+1}$ is measurable, $\psi_i(\eta_i)$ is well-defined. Set $e_i = \psi_i^{-1}(\psi_i(\eta_i))$ and $\omega_i = (e_i, a_i)$. Then, we have that: (i) $\text{marg}_{\mathcal{T}_{-i} \times \Theta} \psi_i(e_i) = \text{marg}_{\mathcal{T}_{-i} \times \Theta} \eta_i$ and (ii) $e_{i,1} = \text{marg}_{A_{-i} \times \Theta} \eta_i$. Thus, it follows that $q_i(e_i) = \tau_i$ and $a_i \in BR_i(e_{i,1})$. Now, notice that for any $k = 0, \dots, n$ we have that,

$$\begin{aligned} \psi_i(e_i(\omega))[\{(\omega'_{-i}, \theta) \in \Omega_{-i} \times \Theta \mid (\omega'_{-i}, \omega_i, \theta) \in R_{-i} \cap B_{-i}^{\lambda,k}(R)\}] &= \\ &= \psi_i(e_i(\omega))[\text{Proj}_{\Omega_{-i}} W_{-i,k} \times \Theta] \\ &= \sum_{\ell=0}^{n+1} \psi_i^\ell(e_i(\omega))[\text{Proj}_{\Omega_{-i}} W_{-i,k} \times \Theta] \\ &\geq \sum_{\ell=k+1}^n \eta_i[\{(\tau_{-i}, a_{-i}) \in (Z_{-i,\ell} \setminus Z_{-i,\ell+1}) \mid (\phi_{-i,\ell}(\tau_{-i}), a_{-i}) \in \text{Proj}_{\Omega_{-i}} W_{-i,k}\} \times \Theta] \\ &\quad + \eta_i[\{(\tau_{-i}, a_{-i}) \in Z_{-i,n+1} \mid (\phi_{-i,n+1}(\tau_{-i}), a_{-i}) \in \text{Proj}_{\Omega_{-i}} W_{-i,k}\} \times \Theta] \\ &\geq \sum_{\ell=k+1}^n \eta_i[(Z_{-i,\ell} \setminus Z_{-i,\ell+1}) \times \Theta] + \eta_i[Z_{-i,n+1} \times \Theta] \\ &= \eta_i[Z_{-i,k+1} \times \Theta] \geq \lambda_{k+1}. \end{aligned}$$

Thus, we conclude that $\omega \in R_i \cap B_i^{\lambda,n+1}(R) \cap [q_i = \tau_i]$, and therefore, that there exists some $(e_i, a_i) \in \text{Proj}_{\mathcal{E}_i \times A_i}(R_i \cap B_i^{\lambda,n+1}(R))$ such that $q_i(e_i) = \tau_i$. Finally, notice that, in particular, the proof of the right inclusion implies that for any $i \in I$ correspondence $\Phi_{i,n+2}$ is non-empty.

Now, in order to finish the proof, fix $i \in I$ and $\tau_i \in \mathcal{T}_i$, and notice that for any

$n \geq 0$,

$$\begin{aligned}
\text{ICR}_{i,n+1}^\lambda(\tau_i) &= \text{Proj}_{A_i}((\{\tau_i\} \times A_i) \cap \text{Graph}(\text{ICR}_{i,n+1}^\lambda)) \\
&= \text{Proj}_{A_i}((\{\tau_i\} \times A_i) \cap \{(q_i(e_i), a_i) \in \mathcal{T}_i \times A_i \mid (e_i, a_i) \in \text{Proj}_{\mathcal{E}_i \times A_i}(R_i \cap B_i^{\lambda,n}(R))\}) \\
&= \text{Proj}_{A_i}(\{\tau_i\} \times \{a_i \in A_i \mid (e_i, a_i) \in \text{Proj}_{\mathcal{E}_i \times A_i}(R_i \cap B_i^{\lambda,n}(R) \cap [q_i = \tau_i])\}) \\
&= \text{Proj}_{A_i}(R_i \cap B_i^{\lambda,n}(R) \cap [q_i = \tau_i]).
\end{aligned}$$

Finally, the fact that,

$$\begin{aligned}
\text{ICR}_i^\lambda(\tau_i) &= \bigcap_{n \geq 0} \text{Proj}_{A_i}(R_i \cap B_i^{\lambda,n}(R) \cap [q_i = \tau_i]) \\
&= \text{Proj}_{A_i}(R_i \cap \bigcap_{n \geq 0} B_i^{\lambda,n}(R) \cap [q_i = \tau_i]) \\
&= \text{Proj}_{A_i}(R_i \cap CB_i^\lambda(R) \cap [q_i = \tau_i]),
\end{aligned}$$

completes the proof. ■

B PROOFS: MAIN RESULTS

B.1 ROBUSTNESS OF THE WY-DISCONTINUITY

PROPOSITION 4 (Robustness of the WY-discontinuity). *Let $\langle \mathcal{G}, \mathcal{T} \rangle$ be a finite Bayesian game satisfying the richness condition. Then, for any player i , any type $t_i \in T_i$ and any $a_i \in \text{ICSR}_i^{\mathcal{T}}(t_i)$ there exists some $p < 1$ and some convergent sequence $(\tau_i^n)_{n \in \mathbb{N}}$ approaching t_i such that $\text{ICR}_i^{\bar{p}, \mathcal{T}}(\tau_i^n) = \{a_i\}$ for any $n \in \mathbb{N}$.*

Proof. Because the type and action spaces are finite, there is a finite N such that for every type t_i of every player, $\text{ICSR}_i(t_i) = \text{ICSR}_{i,n}(t_i)$ for all $n \geq N$. Then there must be, for each $(t_i, a_i) \in \text{Graph}(\text{ICSR}_i)$, a belief $\eta_i[t_i, a_i] \in D_{i,N+1}(t_i)$ with $BR_i(\eta_i[t_i, a_i]) = \{a_i\}$. Then upper-hemicontinuity of BR_i tells us that there is an $\varepsilon[t_i, a_i]$ such that $BR_i(\eta'_i) = \{a_i\}$ for all η'_i with $d(\eta'_i, \eta_i[t_i, a_i]) < \varepsilon[t_i, a_i]$. There are finitely many pairs (t_i, a_i) , so we can let $p < 1$ be such that $1 - p < \varepsilon[t_i, a_i]$ for every (t_i, a_i) . Also, for each a_i , pick a θ_{a_i} as described in the richness assumption and let $p < 1$ be large enough that a_i is the unique best reply for any type who puts probability at least p on θ_{a_i} . We now claim that for such p , the theorem is satisfied.

We will recursively construct a nested sequence of type spaces $\langle T^n, \pi \rangle$, $n = 0, 1, \dots$, with each $T_i^n = \bigcup_{j=0}^n \{[t_i, a_i, j] \mid (t_i, a_i) \in \text{Graph}(\text{ICSR}_i)\}$, such that, for all t_i, a_i , and

j ,

$$\text{ICR}_{i,j+1}^{\bar{p}}([t_i, a_i, j]) = \{a_i\} \quad (1)$$

$$\tau_{i,j}([t_i, a_i, j]) = \tau_{i,j}(t_i), j \geq 1 \quad (2)$$

That is, the belief hierarchy of $[t_i, a_i, j]$ agrees with that of t_i to j orders of belief, and a_i is the only \bar{p} -rationalizable type for each $[t_i, a_i, j]$, which proves the result. Our construction is similar to that in Lemma 7 of [Weinstein and Yildiz \(2007\)](#). Note that we actually achieve a stronger notion of convergence than required: if we used the product topology over discrete topologies on each level of the hierarchy (rather than weak topologies), we would still have convergence. This stronger convergence is possible because of the assumption of *strict* rationalizability.

For $n = 0$, we simply take advantage of the richness assumption: Let $\pi_i([t_i, a_i, 0])([t_{-i}, \theta_{a_i}]) = 1$, where t_{-i} is arbitrary. Property (1) is immediate from the choice of p , and (2) is vacuous in this case.

Now assume we have constructed T^n with the desired properties. We extend the belief map π_i to each T_i^{n+1} by first defining a map

$$\mu_{i,n} : (t_{-i}, a_{-i}, \theta) \mapsto ([t_{-i}, a_{-i}, n], \theta)$$

and then letting

$$\pi_i([t_i, a_i, n+1]) = \eta_i[t_i, a_i] \circ \mu_{i,n}^{-1}$$

Let $S = \{([t_{-i}, a_{-i}, n], a_{-i}, \theta)\} \subseteq T_{-i}^n \times A_{-i} \times \Theta$ be triples where the unique $\text{ICR}_{-i,j+1}^{\bar{p}}$ action a_{-i} of type $[t_{-i}, a_{-i}, n]$ is selected. Now all beliefs $\eta_i \in C_{i,n+1}^{\bar{p}}([t_i, a_i, n+1])$ have marginal $\pi_i([t_i, a_i, n+1])$ on $T_{-i} \times \Theta$ and must satisfy $\eta_i(S) \geq p$. This implies that η_i agrees with $\eta_i[t_i, a_i]$ with probability at least p , implying $d(\eta_i, \eta_i[t_i, a_i]) < \varepsilon[t_i, a_i]$ and $BR_i(\eta_i) = \{a_i\}$, which gives us (1). Furthermore, observe that the map $\mu_{i,n}$ leaves θ unchanged and, by inductive hypothesis, takes (t_{-i}, a_{-i}, θ) to a type $[t_{-i}, a_{-i}, n]$ which agrees with t_{-i} in the first n orders of belief. Also, $\eta_i[t_i, a_i]$ is required to have marginal $\pi_i(t_i)$. These properties of $\mu_{i,n}$ and $\eta_i[t_i, a_i]$ make it immediate that $[t_i, a_i, n+1]$ has the correct order- $n+1$ beliefs, completing the proof. ■

B.2 NON-ROBUSTNESS OF THE WY-DISCONTINUITY

THEOREM 1 (Non-robustness of the WY-discontinuity). *Let \mathcal{G} be a game with incomplete information. Then for any belief hierarchy τ_i and any λ with $\lambda_n \rightarrow 0$ there exists a neighborhood U of τ_i such that $\text{ICSR}_i(\tau_i) \subseteq \text{ICR}_i^\lambda(\hat{\tau}_i)$ for any $\hat{\tau}_i \in U$.*

Proof. We proceed in three steps:

FIRST PRELIMINARY FACT. *For any $i \in I$, any finite $\tau_i \in \mathcal{T}_i$ and any $a_i \in \text{ICSR}_i(\tau_i)$, there exists some $\eta_i \in \bigcap_{n \in \mathbb{N}} D_{i,n}(\tau_i)$ such that $BR_i(\eta_i) = \{a_i\}$.* Fix $i \in I$, finite $\tau_i \in \mathcal{T}_i$ and $a_i \in \text{ICSR}_i(\tau_i)$. Since τ_i is finite, there is some $N \in \mathbb{N}$ for which we can denote $\text{supp}(\text{marg}_{\mathcal{T}_i} \varphi_i(\tau_i)) = \{\tau_{-i}^1, \tau_{-i}^2, \dots, \tau_{-i}^N\}$. Note then that for any $k = 1, \dots, N$ and any $a_{-i} \notin \text{ICSR}_{-i}(\tau_{-i}^k)$ there exists some $m_{a_{-i},k} \in \mathbb{N}$ such that $a_{-i} \in \text{ICSR}_{-i, m_{a_{-i},k}-1}(\tau_{-i}^k) \setminus \text{ICSR}_{-i, m_{a_{-i},k}}(\tau_{-i}^k)$, and thus, since A_{-i} is finite, $m = \max\{m_{a_{-i},k} | k = 1, \dots, N \text{ and } a_{-i} \in A_{-i} \setminus \text{ICSR}_{-i}(\tau_{-i}^k)\}$ is well-defined and satisfies,

$$\text{ICSR}_{-i,m}(\tau_{-i}^k) = \text{ICSR}_{-i}(\tau_{-i}^k), \quad (3)$$

for any $k = 1, \dots, N$. Now, since $a_i \in \text{ICSR}_i(\tau_i)$, in particular, $a_i \in \text{ICSR}_{i,m+1}(\tau_i)$, and therefore, there exists some $\eta_i \in D_{i,m}(\tau_i)$ such that $BR_i(\eta_i) = \{a_i\}$. Obviously, it follows from (3) that $M = \bigcup_{k=1}^N \{\tau_{-i}^k\} \times \text{ICSR}_{-i,m}(\tau_{-i}^k) \subseteq \text{Graph}(\text{ICSR}_{-i,m+\ell})$ for any $\ell \geq 0$, and since it is clearly measurable and we know that $\eta_i[M] = 1$, we conclude that $\eta_i \in \bigcap_{n \in \mathbb{N}} D_{i,n}(\tau_i)$.

SECOND PRELIMINARY FACT. *For any $i \in I$, any $n \geq 0$, any finite $\tau_i \in \mathcal{T}_i$, any $\eta_i \in D_{i,n}(\tau_i)$ and any $\varepsilon > 0$ there exists some $\delta > 0$ such that $C_{i,n}(\hat{\tau}_i) \cap B_\varepsilon(\eta_i) \neq \emptyset$ for any $\hat{\tau}_i \in B_\delta(\tau_i)$.* We proceed by induction. The initial step ($n = 0$) is trivially true, so let's suppose that $n \geq 0$ is such the claim holds.

Note that the induction hypothesis implies that, in particular, the following holds: for any $i \in I$ and any finite $\tau_i \in \mathcal{T}_i$ there exists some open neighborhood of τ_i , $U_{\tau_i}^n$, such that $\text{ICSR}_{i,n+1}(\tau_i) \subseteq \text{ICR}_{i,n+1}(\hat{\tau}_i)$ for any $\hat{\tau}_i \in U_{\tau_i}^{n+1}$. To see this, for each $a_i \in \text{ICSR}_{i,n+1}(\tau_i)$ and $\eta_i \in D_{i,n}(\tau_i)$ such that $BR_i(\eta_i) = \{a_i\}$, pick ε_{η_i} such that $BR_i(\eta'_i) = \{a_i\}$ for any $\eta'_i \in B_{\varepsilon_{\eta_i}}(\eta_i)$ (it exists due to continuity of u_i); then, set $U_{a_i} = B_{\delta_{a_i}}(\tau_i)$ for $\delta_{a_i} > 0$ such that for any $\tau'_i \in B_{\delta_{a_i}}(\tau_i)$ it holds that $C_{i,n}(\tau'_i) \cap B_{\varepsilon_{\eta_i}}(\eta_i) \neq \emptyset$. Obviously, $a_i \in \text{ICR}_{i,n+1}(\tau'_i)$ for any $\tau'_i \in U_{a_i}$, and hence, $\text{ICSR}_{i,n+1}(\tau_i) \subseteq \text{ICR}_{i,n+1}(\hat{\tau}_i)$ for any $\hat{\tau}_i \in U_{\tau_i}^{n+1} = \bigcap_{a_i \in \text{ICSR}_{i,n+1}(\tau_i)} U_{a_i}$.

Now, fix $i \in I$, finite $\tau_i \in \mathcal{T}_i$ and $\eta_i \in D_{i,n+1}(\tau_i)$. We know from the induction hypothesis that for any $\tau_{-i} \in \text{supp}(\text{marg}_{\mathcal{T}_i} \varphi_i(\tau_i))$ there exists some $\delta_{\tau_{-i}} > 0$ such that $\text{ICSR}_{-i,n+1}(\tau_{-i}) \subseteq \text{ICR}_{-i,n+1}(\tau'_{-i})$ for any $\tau'_{-i} \in B_{\delta_{\tau_{-i}}}(\tau_{-i})$; furthermore, since τ_i is finite and \mathcal{T}_i is metrizable, we can assume w.l.o.g. that there exists some $\delta^* > 0$ such that (i) $B_{\delta^*}(\tau_{-i}) \subseteq B_{\delta_{\tau_{-i}}}(\tau_{-i})$ for any $\tau_{-i} \in \text{supp}(\text{marg}_{\mathcal{T}_i} \varphi_i(\tau_i))$ and (ii) $B_{2\delta^*}(\tau_{-i}, \theta) \cap B_{2\delta^*}(\tau'_{-i}, \theta') = \emptyset$ for every distinct $(\tau_{-i}, \theta), (\tau'_{-i}, \theta') \in \text{supp} \varphi_i(\tau_i)$.

Next, for any $\delta' \leq \delta^*$ denote $S = \text{supp} \varphi_i(\tau_i)$, and let $S^{\delta'} = \bigcup_{(\tau_{-i}, \theta) \in S} B_{\delta'}(\tau_{-i}, \theta)$. Also, for each $(\tau_{-i}, \theta) \in S^{\delta'}$, let $c(\tau_{-i}, \theta)$ be the unique element in S such that

$(\tau_{-i}, \theta) \in B_{\delta'}(c(\tau_{-i}, \theta))$ (where uniqueness is assured by the definition of δ^*).²² Then, for any $\tau_i^{\delta'} \in \varphi_i^{-1}(B_{\delta'}(\varphi_i(\tau_i)))$ we define conjecture $\eta_i^{\delta'} \in C_{i,n+1}(\tau_i^{\delta'})$ by picking arbitrary $\bar{\eta}_i \in C_{i,n+1}(\tau_i^{\delta'})$ and setting:

$$\eta_i^{\delta'}[(\tau_{-i}, a_{-i}, \theta)] = \begin{cases} \varphi_i(\tau_i^{\delta'})[(\tau_{-i}, \theta)] \cdot \left(\frac{\eta_i[\{c(\tau_{-i}, \theta)\} \times \{a_{-i}\}]}{\varphi_i(\tau_i)[c(\tau_{-i}, \theta)]} \right) & \text{if } (\tau_{-i}, \theta) \in S^{\delta'}, \\ \bar{\eta}_i[(\tau_{-i}, a_{-i}, \theta)] & \text{if } (\tau_{-i}, \theta) \notin S^{\delta'} \text{ and } a_{-i} \in \text{ICR}_{-i,n+1}(\tau_{-i}), \\ 0 & \text{otherwise,} \end{cases}$$

for any $(\tau_{-i}, a_{-i}, \theta) \in \mathcal{T}_{-i} \times A_{-i} \times \Theta$. Finiteness of τ_i and A_{-i} and disjointness of family of balls $\{B_{\delta'}(\tau_{-i}, \theta) | (\tau_{-i}, \theta) \in S\}$ guarantee that $\eta_i^{\delta'}$ is indeed a well-defined element of $\Delta(\mathcal{T}_{-i} \times A_{-i} \times \Theta)$. Note in addition that the following hold:

$$(i) \text{ marg}_{\mathcal{T}_{-i} \times \Theta} \eta_i^{\delta'} = \varphi_i(\tau_i^{\delta'}).$$

$$(ii) \eta_i^{\delta'}[\text{Graph}(\text{ICR}_{-i,n+1}) \times \Theta] = 1. \text{ To see this, notice that for any } (\tau_{-i}, \theta) \in \text{supp } \varphi_i(\tau_i^{\delta'}) \cap S^{\delta'}, \text{ if we denote } (\tau'_{-i}, \theta') = c(\tau_{-i}, \theta),$$

$$\begin{aligned} \eta_i^{\delta'}[\{(\tau_{-i}, \theta)\} \times \text{ICR}_{-i,n+1}(\tau_{-i})] &= \varphi_i(\tau_i^{\delta'})[(\tau_{-i}, \theta)] \cdot \left(\frac{\eta_i[\{c(\tau_{-i}, \theta)\} \times \text{ICR}_{-i,n+1}(\tau_{-i})]}{\varphi_i(\tau_i)[c(\tau_{-i}, \theta)]} \right) \\ &= \varphi_i(\tau_i^{\delta'})[(\tau_{-i}, \theta)] \cdot \left(\frac{\eta_i[\{c(\tau_{-i}, \theta)\} \times \text{ICSR}_{-i,n+1}(\tau'_{-i})]}{\varphi_i(\tau_i)[c(\tau_{-i}, \theta)]} \right) \\ &= \varphi_i(\tau_i^{\delta'})[(\tau_{-i}, \theta)], \end{aligned}$$

Note that in the second equality we make use of the property derived from the induction hypothesis at the beginning of the proof.²³ Remember that due to the fact that $\bar{\eta}_i \in C_{i,n+1}(\tau_i^{\delta'})$, and thus, for any $(\tau_{-i}, \theta) \in \text{supp } \varphi_i(\tau_i^{\delta'}) \cap (S^{\delta'})^c$, we have that:

$$\eta_i^{\delta'}[\{(\tau_{-i}, \theta)\} \times \text{ICR}_{-i,n+1}(\tau_{-i})] = \varphi_i(\tau_i^{\delta'})[(\tau_{-i}, \theta)].$$

We see then that $\eta_i^{\delta'} \in C_{i,n+1}(\tau_i^{\delta'})$. We are going to check next that as δ' tends to zero, $d(\eta_i, \eta_i^{\delta'})$, where d denotes the Lévy-Prohorov metric, gets arbitrarily small. Specifically, we are going to check that $d(\eta_i, \eta_i^{\delta'}) \leq (|S| + 1)\delta'$ for every $\delta' \leq \delta^*$. To see this, we need to check that:

²²That is, disjointness of family $\{B_{\delta'}(\tau_{-i}, \theta) | (\tau_{-i}, \theta) \in S\}$ implies that for each $(\tau_{-i}, \theta) \in \mathcal{T}_{-i} \times \Theta$, there exists at most one $B \in \{B_{\delta'}(\tau_{-i}, \theta) | (\tau_{-i}, \theta) \in S\}$ such that $(\tau_{-i}, \theta) \in B$. Let $c(\tau_{-i}, \theta)$ denote the center of B .

²³I.e., it follows from the facts that $\text{ICSR}_{-i,n+1}(\tau'_{-i}) \subseteq \text{ICR}_{-i,n+1}(\tau_{-i})$ and $\eta_i \in D_{i,n+1}(\tau_i)$.

- (i) $\eta_i^{\delta'}[E] \leq \eta_i[E^{(|S|+1)\delta'}] + (|S|+1)\delta'$ for any measurable E .
- (ii) $\eta_i[E] \leq \eta_i^{\delta'}[E^{(|S|+1)\delta'}] + (|S|+1)\delta'$ for any measurable E .

Before checking (i) and (ii) let us make a remark and introduce some notation:

- Remember that $\varphi_i(\tau_i^{\delta'}) \in B_{\delta'}(\varphi_i(\tau_i))$; in particular this implies that for any $(\tau_{-i}, \theta) \in S$ we have that:

$$\varphi_i(\tau_i)[(\tau_{-i}, \theta)] + \delta' = \varphi_i(\tau_i)[B_{2\delta'}(\tau_{-i}, \theta)] + \delta' \geq \varphi_i(\tau_i)[(B_{\delta'}(\tau_{-i}, \theta))^{\delta'}] + \delta' \geq \varphi_i(\tau_i^{\delta'})[B_{\delta'}(\tau_{-i}, \theta)].$$

- For each measurable E denote $X(E) = \eta_i^{\delta'}[E \cap S^{\delta'}]$ and $Y(E) = \eta_i^{\delta'}[E \cap (S^{\delta'})^c]$. Obviously, $\eta_i^{\delta'}[E] = X(E) + Y(E)$ and in addition, it is easy to check that:

$$X(E) = \sum_{(\tau_{-i}, \theta) \in S} \left(\frac{\varphi_i(\tau_i^{\delta'})[B_{\delta'}(\tau_{-i}, \theta)]}{\varphi_i(\tau_i)[(\tau_{-i}, \theta)]} \right) \cdot \eta_i[\{(\tau_{-i}, \theta)\} \times \{a_{-i} \in A_{-i} | (\tau_{-i}, a_{-i}, \theta) \in E\}],$$

$$Y(E) = \bar{\eta}_i[E \cap (S^{\delta'})^c].$$

Now, let's prove (i) and (ii) above:

- (i) To see this, notice first that:

$$\begin{aligned} X(E) &= \sum_{(\tau_{-i}, \theta) \in S} \left(\frac{\varphi_i(\tau_i^{\delta'})[B_{\delta'}(\tau_{-i}, \theta)]}{\varphi_i(\tau_i)[(\tau_{-i}, \theta)]} \right) \cdot \eta_i[\{(\tau_{-i}, \theta)\} \times \{a_{-i} \in A_{-i} | (\tau_{-i}, a_{-i}, \theta) \in E\}] \\ &\leq \sum_{(\tau_{-i}, \theta) \in S} \left(\frac{\varphi_i(\tau_i)[(\tau_{-i}, \theta)] + \delta'}{\varphi_i(\tau_i)[(\tau_{-i}, \theta)]} \right) \cdot \eta_i[\{(\tau_{-i}, \theta)\} \times \{a_{-i} \in A_{-i} | (\tau_{-i}, a_{-i}, \theta) \in E\}] \\ &= \sum_{(\tau_{-i}, \theta) \in S} \left(1 + \frac{\delta'}{\varphi_i(\tau_i)[(\tau_{-i}, \theta)]} \right) \cdot \eta_i[\{(\tau_{-i}, \theta)\} \times \{a_{-i} \in A_{-i} | (\tau_{-i}, a_{-i}, \theta) \in E\}] \\ &\leq \eta_i[E^{\delta'}] + |S|\delta'. \end{aligned}$$

In addition,

$$Y(E) \leq \bar{\eta}_i[A_{-i} \times \text{Proj}_{\mathcal{T}_{-i} \times \Theta}(E) \cap (S^{\delta'})^c] \leq \varphi_i(\tau_i^{\delta'})[(S^{\delta'})^c] \leq \varphi_i(\tau_i)\left[\left((S^{\delta'})^c\right)^{\delta'}\right] + \delta' = \delta'.$$

Thus, the claim in (i) follows immediately.

(ii) To see this, notice first that:

$$\begin{aligned}
X(E) &= \sum_{(\tau_{-i}, \theta) \in S} \left(\frac{\varphi_i(\tau_i^{\delta'})[B_{\delta'}(\tau_{-i}, \theta)]}{\varphi_i(\tau_i)[(\tau_{-i}, \theta)]} \right) \cdot \eta_i [\{(\tau_{-i}, \theta)\} \times \{a_{-i} \in A_{-i} | (\tau_{-i}, a_{-i}, \theta) \in E\}] \\
&\geq \sum_{(\tau_{-i}, \theta) \in S} \left(\frac{\varphi_i(\tau_i^{\delta'})[B_{\delta'}(\tau_{-i}, \theta)]}{\varphi_i(\tau_i^{\delta'})[B_{\delta'}(\tau_{-i}, \theta)] + \delta'} \right) \cdot \eta_i [\{(\tau_{-i}, \theta)\} \times \{a_{-i} \in A_{-i} | (\tau_{-i}, a_{-i}, \theta) \in E\}] \\
&= \sum_{(\tau_{-i}, \theta) \in S} \left(1 - \frac{\delta'}{\varphi_i(\tau_i^{\delta'})[B_{\delta'}(\tau_{-i}, \theta)] + \delta'} \right) \cdot \eta_i [\{(\tau_{-i}, \theta)\} \times \{a_{-i} \in A_{-i} | (\tau_{-i}, a_{-i}, \theta) \in E\}] \\
&\geq \eta_i[E] - |S|\delta'.
\end{aligned}$$

Then, (ii) follows since:

$$\begin{aligned}
\eta_i^{\delta'}[E^{(|S|+1)\delta'}] + (|S|+1)\delta' &\geq \eta_i^{\delta'}[E] + (|S|+1)\delta' \geq \\
&\geq X(E) + (|S|+1)\delta' \geq \eta_i[E] - |S|\delta' + (|S|+1)\delta' \geq \eta_i[E].
\end{aligned}$$

We conclude then that that $d(\eta_i, \eta_i^{\delta'}) \leq (|S|+1)\delta'$ for every $\delta' \leq \delta^*$. Consequently, for any $\varepsilon > 0$ and any $\delta(\varepsilon) \leq \min\{\delta^*, \varepsilon/(|S|+1)\}$ we have that $d(\eta_i, \eta_i^{\delta(\varepsilon)}) < \varepsilon$. Now, finally, the fact that $\varphi_i^{-1}(B_{\delta(\varepsilon)}(\varphi_i(\tau_i)))$ is open implies that there exists some $\delta > 0$ such that $B_\delta(\tau_i) \subseteq \varphi_i^{-1}(B_{\delta(\varepsilon)}(\varphi_i(\tau_i)))$ and hence it follows that $C_{i,n+1}(\hat{\tau}_i) \cap B_\varepsilon(\eta_i) \neq \emptyset$ for any $\hat{\tau}_i \in B_\delta(\tau_i)$.

PROOF OF THE THEOREM. Fix $i \in I$, finite $\tau_i \in \mathcal{T}_i$ and λ with $\lambda_n \rightarrow 0$. Pick $a_i \in \text{ICSR}_i(\tau_i)$ and $\eta_i \in \bigcap_{n \geq 0} D_{i,n}(\tau_i)$ such that $BR_i(\eta_i) = \{a_i\}$ (we know from the first preliminary fact that such η_i exists). Then, it follows from continuity of u_i that:

(i) There exists some $\bar{\varepsilon} > 0$ such that $BR_i(\eta'_i) = \{a_i\}$ for any $\eta'_i \in B_{\bar{\varepsilon}}(\eta_i)$.

From the second preliminary fact we know that:

(ii) For any $n \in \mathbb{N}$ there exists some $\delta_{a_i} > 0$ such that for any $\tau'_i \in B_{\delta_{a_i}}(\tau_i)$ there exists some $\eta'_i \in C_{i,n}(\tau'_i) \cap B_{\bar{\varepsilon}/2}(\eta_i)$.

Because $\lambda_n \rightarrow 0$ we also have:

(iii) There exists some $n_{a_i} \in \mathbb{N}$ such that $\lambda_k < \bar{\varepsilon}/2$ for $k > n_{a_i}$.

Thus, it follows from (i)–(iii) that for each a_i there exists some $\delta_{a_i} > 0$ such that, for any $\tau'_i \in B_{\delta_{a_i}}(\tau_i)$, any $k > n_{a_i}$ and any $\eta'_i \in C_{i,n_{a_i}}(\tau'_i) \cap B_{\bar{\varepsilon}/2}(\eta_i) \neq \emptyset$ we have that,

$$d((1 - \lambda_k) \cdot \eta'_i + \lambda_k \cdot \mu_i, \eta_i) < d(\eta'_i, \eta_i) + \lambda_k < \bar{\varepsilon},$$

for any $\mu_i \in \Delta(\mathcal{T}_{-i} \times A_{-i} \times \Theta)$. Now, fix $\hat{\tau}_i \in B_{\delta_{a_i}}(\tau_i)$, fix arbitrary conjectures $\eta'_i \in C_{i,n_{a_i}}(\hat{\tau}_i) \cap B_{\bar{\varepsilon}/2}(\eta_i)$ and $\bar{\eta}_i \in \bigcap_{n \in \mathbb{N}} C_{i,n}(\hat{\tau}_i)$, and define:

$$\hat{\eta}_i = \left(1 - \frac{\bar{\varepsilon}}{2}\right) \cdot \eta'_i + \left(\frac{\bar{\varepsilon}}{2}\right) \cdot \bar{\eta}_i.$$

Obviously, $\text{marg}_{\mathcal{T}_{-i} \times \Theta} \hat{\eta}_i = \varphi_i(\hat{\tau}_i)$, and since $\hat{\eta}_i \in B_{\bar{\varepsilon}}(\eta_i)$, we have that $BR_i(\hat{\eta}_i) = \{a_i\}$. Now notice first that, for any $n \leq n_{a_i}$,

$$\hat{\eta}_i [\Theta \times \text{Graph}(\text{ICR}_{-i,n}^\lambda)] \geq \hat{\eta}_i [\Theta \times \text{Graph}(\text{ICR}_{-i,n})] = \left(1 - \frac{\bar{\varepsilon}}{2}\right) + \frac{\bar{\varepsilon}}{2} \geq \lambda_n,$$

and hence $\hat{\eta}_i \in C_{i,n}^\lambda(\hat{\tau}_i)$. Note in addition that, for any $n > n_{a_i}$,

$$\begin{aligned} \hat{\eta}_i [\Theta \times \text{Graph}(\text{ICR}_{-i,n}^\lambda)] &\geq \left(\frac{\bar{\varepsilon}}{2}\right) \cdot \bar{\eta}_i [\Theta \times \text{Graph}(\text{ICR}_{-i,n}^\lambda)] \\ &\geq \left(\frac{\bar{\varepsilon}}{2}\right) \cdot \bar{\eta}_i [\Theta \times \text{Graph}(\text{ICR}_{-i,n})] \\ &= \frac{\bar{\varepsilon}}{2} \geq \lambda_n. \end{aligned}$$

Thus, for any $n \geq 0$ we have $\hat{\eta}_i \in C_{i,n}^\lambda(\hat{\tau}_i)$ and consequently $a_i \in \text{ICR}_{i,n+1}^\lambda(\hat{\tau}_i)$. Hence, we conclude that $a_i \in \text{ICR}_i^\lambda(\hat{\tau}_i)$ for any $\hat{\tau}_i \in B_{\delta_{a_i}}(\tau_i)$. It follows that $\text{ICSR}_i(\tau_i) \subseteq \text{ICR}_i^\lambda(\hat{\tau}_i)$ for any $\hat{\tau}_i \in U = \bigcap_{a_i \in \text{ICSR}_i(\tau_i)} B_{\delta_{a_i}}(\tau_i)$. \blacksquare

COROLLARY 1 (Non-robustness of generic uniqueness). *Let \mathcal{G} be a game with incomplete information and finite set of action profiles. Then, for any player i for which there exists some τ_i such that $|\text{ICSR}_i(\tau_i)| > 1$, and for any λ with $\lambda_n \rightarrow 0$, the following set is not dense:*

$$\mathcal{U}_i^\lambda = \{\tau_i \in \mathcal{T}_i \mid |\text{ICR}_i^\lambda(\tau_i)| = 1\}.$$

Proof. It follows directly from Proposition 1: fix sequence λ with limit 0 and pick $i \in I$ and $\tau_i \in \mathcal{T}_i$ such that $|\text{ICSR}_i(\tau_i)| > 1$. Then, we know that there exists some open neighborhood U of τ_i such that $\text{ICSR}_i(\tau_i) \not\subseteq \text{ICR}_i^\lambda(\hat{\tau}_i)$ for any $\hat{\tau}_i \in U$. Thus, $\mathcal{U}_i^\lambda \subseteq \mathcal{T}_i \setminus U$, and hence, it is not dense. \blacksquare

REFERENCES

Alaoui, Larbi and Antonio Penta (2016). “Endogenous depth of reasoning”. *The Review of Economic Studies*, **83**, 1297–1333.

- Aliprantis, Charambolos D. and Kim C. Border (1999). *Infinite dimensional analysis: A hitchhiker's guide, 3rd edition*. Springer, Berlin.
- Ambrosio, Luigi, Nicola Gigli and Giuseppe Savaré (2006). *Gradient flows in metric spaces and in the space of probability measures*. Birkhauser Verlag.
- 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.
- Brandenburger, Adam and Eddie Dekel (1993). “Hierarchies of beliefs and common knowledge”. *Journal of Economic Theory*, **59**, 189–198.
- Camerer, Colin F., George Loewenstein and Matthew Rabin (editors) (2004). *Advances in behavioral economics*. New York, Russel, Sage Foundation.
- Carlsson, Hans and Eric van Damme (1993). “Global games and equilibrium selection”. *Econometrica*, **61**, 989–1018.
- 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.
- Chen, Yi-Chun, Xiao Luo and Chen Qu (2015). “Rationalizability in general situations”. *Economic Theory*, pp. 1–21.
- Chen, Yi-Chun, Satoru Takahashi and Siyan (2014). “The Weinstein-Yildiz Selection and Robust Predictions with Arbitrary Payoff Uncertainty”.
- Chen, Yi-Chun, Satoru Takahashi and Siyang Xiong (2014). “The robust selection of rationalizability”. *Journal of Economic Theory*, **151**, 448–475.
- Chen, Yi-Chun, Alfredo Di Tillio, Eduardo Faingold and Siyang Xiong (2016). “Characterizing the stratategic impact of misspecified beliefs”. *The Review of Economic Studies*. (Forthcoming).
- 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.

- Friedenberg, Amanda, Willemien Kets and Terri Kneeland (2016). “Bounded reasoning: rationality or cognition?”. Mimeo.
- Germano, Fabrizio and Peio Zuazo-Garin (2017). “Bounded rationality and correlated equilibria”. *International Journal of Game Theory*, **vv**, xx–yy.
- 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 Willemien Kets (2016). “Robust multiplicity with a grain of naiveté”. Mimeo.
- 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 equilibria to incomplete information”. *Econometrica*, **65**, 1283–1309.
- Kets, Willemien (2014). “Finite depth of reasoning and equilibrium play”. Mimeo.
- Monderer, Dov and Dov Samet (1987). “Approximating common knowledge with common beliefs”. *Games and Economic Behavior*, **1**, 170–190.
- Monderer, Dov and Dov Samet (1996). “Proximity of information in games with incomplete information”. *Mathematics of Operations Research*, **21**, 707–725.
- Morris, Stephen and Hyun Song Shin (1998). “Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks”. *American Economic Review*, **88**, 587–597.
- Nagel, Rosemarie (1995). “Unraveling the guessing game: an experimental study”. *American Economic Review*, **83**, 1313–1326.
- 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.
- Rubinstein, Ariel (1989). “The electronic mail game: strategic behaviour under ‘almost common knowledge’”. *American Economic Review*, **79**, 385–391.

- Stahl, Dale O. and Paul W. Wilson (1994). “Experimental evidence on players’ models of other players”. *Journal of Economic Behavior & Organization*, **25**, 309–327.
- Strzalecki, Tomasz (2010). “Depth of reasoning and higher order beliefs”. Mimeo.
- Weinstein, Jonathan and Muhamet Yildiz (2007). “A structure theorem for rationalizability with application to robust predictions of refinements”. *Econometrica*, **75**, 365–400.