Theoretical Economics 15 (2020), 89–122 1555-7561/20200089

Uncertain rationality, depth of reasoning and robustness in games with incomplete information

Fabrizio Germano
Department of Economics and Business, Universitat Pompeu Fabra and Barcelona GSE

Jonathan Weinstein
Department of Economics, Washington University in St. Louis

Peio Zuazo-Garin
Department of Foundations of Economic Analysis I, University of the Basque Country

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$^{\lambda}$, where λ is a sequence whose term λ_n is the probability players attach to $(n-1)$th-order belief in rationality. We find that Weinstein and Yildiz's discontinuity remains when λ_n is above an appropriate threshold for all n, but fails when λ_n converges to 0. That is, if players' confidence in mutual rationality persists at high orders, the discontinuity persists, but if confidence vanishes at high orders, the discontinuity vanishes.

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

JEL classification. C72, D82, D83.

Fabrizio Germano: fabrizio.germano@upf.edu
Jonathan Weinstein: j.weinstein@wustl.edu
Peio Zuazo-Garin: peio.zuazo@ehu.eus

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, CESC1 and GAMES 2016. Germano acknowledges financial support from the Spanish Ministry of Economy and Competitiveness (grant ECO2017-89240-P, AEI/FEDER, UE, and through the Severo Ochoa Program for Centers of Excellence in R&D, grant SEV-2015-0563). Zuazo-Garin acknowledges financial help from the Spanish Ministry of Economy and Competitiveness, 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 the ERC Starting Grant #759424, and thanks Northwestern University (both the Department of Economics and the Math Center at MEDS Department) for generous hospitality.

Available at http://econtheory.org. https://doi.org/10.3982/TE2734

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 that would be rationalizable under full common knowledge will be nonrationalizable 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 that 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 et al. (2010) 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 assumes 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 nth-order belief in rationality; ICR itself would be the special case that $\lambda = (1, 1, \dots)$. The answer is twofold: when each element λ_n is above a threshold close enough to 1, we find that WY-discontinuity remains (Proposition 5), but when $(\lambda_n)_{n\in\mathbb{N}} \to 0$ as $n \to \infty$ we find that continuity is restored (Proposition 6). 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 nontrivial refinements are nonrobust 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 5 and 6, 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, i.e., 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, ICR^λ, as a correspondence of 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 7). In particular, all these results, besides Propositions 5 and 6, 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	θ, θ	$\theta-1$, 0
No attack	0, $\theta-1$	0, 0

for $\theta \in \left\{-\frac{2}{5}, \frac{2}{5}\right\}$.

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 that 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 5: 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 $\lambda_k \to 0$, 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 6: 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 (2014) and Heifetz and Kets (2018) 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 (2018) 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 rationalizability considered replacing the product topology with alternate notions of proximity. Dekel et al. (2006) introduce the *strategic topology* that 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 et al. (2010, 2017) 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 et al. 2014a, 2014b).

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, 1968a, 1968b) 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 finite set of *payoff states*, and for each player i we have (iii) a finite set of *actions*, A_i, and (iv) a *utility function* $u_i : A \times \Theta \to [-M, M]$, where $A = \prod_{i\in I} A_i$ denotes the set of action *profiles*.[2] For each player i, we refer to a

[1]The connection of this paper and Heifetz and Kets (2018) is examined in further detail in Section 4.2.2.

[2]In a previous version, we allowed Θ to be any compact metric space, and our main results still hold in that case. We switched to the finite case because it connects more closely with the previous literature and so significantly shortens our proofs.

probability measure $\mu_i \in \Delta(A_{-i} \times \Theta)$,[3] where $A_{-i} = \prod_{j \neq i} A_j$, as a *conjecture*, and for each $\varepsilon \in \mathbb{R}$ player i's *ε-best-reply correspondence* is $\varepsilon\text{-}BR_i : \Delta(A_{-i} \times \Theta) \rightrightarrows A_i$, given by

$$\mu_i \mapsto \left\{ a_i \in A_i \;\middle|\; \begin{array}{l} \text{For every } a_i' \in A_i \setminus \{a_i\}, \\ \displaystyle\int_{A_{-i} \times \Theta} \big(u_i\big((a_{-i}; a_i), \theta\big) - u_i\big((a_{-i}; a_i'), \theta\big)\big) \, \mathrm{d}\mu_i \geq -\varepsilon \end{array} \right\}.$$

When $\varepsilon = 0$, the ε-best-reply correspondence boils down to the standard best-reply correspondence and in such case, we simply denote it by BR_i. Notice that due to the topological assumptions specified above, the ε-best-reply correspondence is known to be nonempty when $\varepsilon \geq 0$, and upper-hemicontinuous for all $\varepsilon \in \mathbb{R}$.[4] We typically represent players' beliefs over Θ by endowing $\mathcal{G}$ with a *type structure* à la Harsanyi (1967). 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* $\pi_i : T_i \to \Delta(T_{-i} \times \Theta)$ 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 *nth-order belief*, and to the elements of $\mathcal{T}_i^0 = \prod_{n \in \mathbb{N}} Z_i^n$, as *belief hierarchies*. A belief hierarchy τ_i is said to be *coherent* if higher-order beliefs do not contradict lower order ones, i.e., if $\mathrm{marg}_{X_i^n} \tau_{i,n+1} = \tau_{i,n}$ 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.[5] Brandenburger and Dekel (1993) show that there exists a homeomorphism $\varphi_i : \mathcal{T}_i \to \Delta(\mathcal{T}_{-i} \times \Theta)$, with $\mathcal{T}_{-i} = \prod_{j \neq i} \mathcal{T}_j$, such that $\mathrm{marg}_{X_i^n} \varphi_i(\tau_i) = \tau_{i,n}$ for any belief hierarchy τ_i and any $n \in \mathbb{N}$. Obviously, $\mathcal{T}^* = \langle \mathcal{T}_i, \varphi_i \rangle_{i \in I}$ 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

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

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

[5] Formally, $\mathcal{T}_i = \bigcap_{n \geq 0} \mathcal{T}_i^n$, where $\mathcal{T}_i^{n+1} = \{\tau_i \in \mathcal{T}_i^n | \tau_{i,m}[\mathrm{Proj}_{X_i^m}(\mathcal{T}_{-i}^n \times \Theta)] = 1 \text{ for any } m \in \mathbb{N}\}$ 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 $\mathrm{Proj}_X S = \{x \in X | (x, y) \in S \text{ for some } y \in Y\}$.

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, 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 $\tau_{i,1}(t_i) = \operatorname{marg}_\Theta \pi_i(t_i)$ and then, for any $n \in \mathbb{N}$ define $(n+1)$th-order belief $\tau_{i,n+1}(t_i)$ by setting,

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

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 $\tau_{i,n} : T_i \to Z_i^n$ is continuous. In addition, is is easy to see that $\tau_i(T_i) \subseteq \mathcal{T}_i$; thus, $\tau_i : T_i \to \mathcal{T}_i$ is a well-defined continuous map. Furthermore, if T_i has nonredundant types,[6] then it is homeomorphic to $\tau_i(T_i)$.

2.3 *Rationalizability and the WY-discontinuity*

Once a player's uncertainty with respect to 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), introduced by Dekel et al. (2007). First, let us recall the more general version of ICR embodied by ε-ICR due to Dekel et al. (2006). Given a Bayesian game $\langle \mathcal{G}, \mathcal{T} \rangle$ and a real number ε, player i's set of (*interim correlated*) *ε-rationalizable* (ε-ICR) actions for type t_i is defined as $\varepsilon\text{-ICR}_i(t_i) = \bigcap_{n\geq 0} \varepsilon\text{-ICR}_{i,n}(t_i)$, where

$$\varepsilon\text{-ICR}_{i,0}(t_i) = A_i,$$
$$\varepsilon\text{-C}_{i,0}(t_i) = \big\{\mu_i \in \Delta(T_{-i} \times A_{-i} \times \Theta) |\operatorname{marg}_{T_{-i}\times\Theta} \mu_i = \pi_i(t_i)\big\},$$

and recursively,[7]

$$\varepsilon\text{-ICR}_{i,n}(t_i) = \big\{a_i \in A_i | a_i \in \varepsilon\text{-}BR_i(\mu_i) \text{ for some } \mu_i \in \varepsilon\text{-C}_{i,n-1}(t_i)\big\},$$
$$\varepsilon\text{-C}_{i,n}(t_i) = \big\{\mu_i \in \varepsilon\text{-C}_{i,0}(t_i) | \operatorname{supp} \mu_i \subseteq \operatorname{Graph}(\varepsilon\text{-ICR}_{-i,n}) \times \Theta\big\},$$

for any $n \in \mathbb{N}$. In the case of $\varepsilon = 0$, the definition collapses to Dekel et al.'s (2007) (*interim correlated*) *rationalizability* (ICR) and in such case we denote the resulting correspondence simply by ICR_i. Dekel et al. (2007) and Battigalli et al. (2011) show that when

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

[7]In addition, let us denote $\operatorname{Graph}(\varepsilon\text{-ICR}_{-i,n}) = \prod_{j\neq i}\{(t_j, a_j) \in T_j \times A_j | a_j \in \varepsilon\text{-ICR}_{j,n}(t_j)\}$.

$\varepsilon = 0$ the specific type structure employed to codify belief hierarchies is immaterial: 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 et al. (2006) that the correspondence $\varepsilon\text{-ICR}_i : T_i \rightrightarrows A_i$ is upper-hemicontinuous, and by Dekel et al. (2007) and Battigalli et al. (2011), that ICR characterizes the behavioral implications of rationality and common belief in rationality. Notice that we permit both positive and negative ε; when $\varepsilon > 0$, the concept is more permissive than standard ICR, and when $\varepsilon < 0$ it is more strict.

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. To better understand this phenomenon, let us recall the *richness condition* first.

Definition 1 (Richness condition). We say that a Bayesian game 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 $a_i' \neq a_i$.

That is, in games that satisfy the richness condition, no action is commonly known not to be strictly dominant. In this context, the main result by Weinstein and Yildiz (2007) tells us that for any type t_i and any action $a_i \in \text{ICR}_i(t_i)$ there exists some sequence of belief hierarchies $(\tau_i^n)_{n\in\mathbb{N}}$ converging to $\tau_i(t_i)$ such that $\text{ICR}_i(\tau_i^n) = \{a_i\}$ for any $n \in \mathbb{N}$.[9] This property, which we refer to as the *WY-discontinuity*, has important implications for games with incomplete information:

- *Nonrobustness of refinements*. No nontrivial 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 nontrivial 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 $\varnothing \neq S_i(\tau_i^n) \subseteq \text{ICR}_i(\tau_i^n) = \{a_i\}$ 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.

[8] That is, for any player i and any type t_i it holds that $\text{ICR}_i(t_i) = \text{ICR}_i(\tau_i(t_i))$.

[9] Recently, 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.

3. Interim correlated λ-rationalizability

3.1 *Definition*

We now introduce interim correlated λ-rationalizability (ICR^λ), the solution concept that formalizes 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 performing stage n of the elimination process, players have confidence λ_n that others have followed the elimination process at previous stages, as captured in the following definition.

Definition 2 (Interim correlated λ-rationalizability). 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 $\mathrm{ICR}_i^\lambda(t_i) = \bigcap_{n\geq 0} \mathrm{ICR}_{i,n}^\lambda(t_i)$, where

$$\mathrm{ICR}_{i,0}^\lambda(t_i) = A_i,$$

$$\mathrm{C}_{i,0}^\lambda(t_i) = \big\{\mu_i \in \Delta(T_{-i} \times A_{-i} \times \Theta) | \operatorname{marg}_{T_{-i}\times\Theta} \mu_i = \pi_i(t_i)\big\},$$

and recursively, for any $n \in \mathbb{N}$,[10]

$$\mathrm{ICR}_{i,n}^\lambda(t_i) = \big\{a_i \in \mathrm{ICR}_{i,n-1}^\lambda(t_i) | a_i \in BR_i(\mu_i) \text{ for some } \mu_i \in \mathrm{C}_{i,n-1}^\lambda(t_i)\big\},$$

$$\mathrm{C}_{i,n}^\lambda(t_i) = \big\{\mu_i \in \mathrm{C}_{i,n-1}^\lambda(t_i) | \mu_i\big[\mathrm{Graph}\big(\mathrm{ICR}_{-i,n}^\lambda\big) \times \Theta\big] \geq \lambda_n\big\}.$$

For $p \in [0,1]$, we will use $\lambda = \bar{p}$ to signify the constant sequence $\lambda_n \equiv p$. Then $\mathrm{ICR}^{\bar{p}}$ will reflect reasoning that is depth-independent, capturing departures from common belief in rationality in the sense of (A), but not (B) above. Also, we use the usual termwise partial ordering on sequences, so in particular $\lambda \geq \bar{p}$ will mean that $\lambda_n \geq p$ for all n. Our examples all focus on the natural case of decreasing λ, though we do not require this in the definition. This case is natural because it represents depth-dependent reasoning that is less confident at higher orders, hence capturing both (A) and (B) above. We will especially consider the case $\lambda \to 0$, 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 (1989) 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 $\mathrm{ICR}^{\bar{p}}$ actions as *interim correlated p-rationalizable.*

[10] We will show later, in Remark 1, that the set $\mathrm{Graph}(\mathrm{ICR}_{-i,n}^\lambda)$ appearing in the last equation of this definition is indeed always measurable.

(ii) *Rationalizability.* The special case $\lambda = \bar{1}$. The standard case of common belief in rationality, i.e., infinite depth of reasoning in which player adhere probability 1 to rationality at every iteration. The case $\mathrm{ICR}^{\bar{1}}$ reduces to the standard notion of ICR, as defined by Dekel et al. (2007) and discussed above.

(iii) *Models with k orders of belief in rationality.* For each $k \geq 0$, define sequence $\lambda^k = (\lambda^k_n)_{n\in\mathbb{N}}$ by

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

An action $a_i \in \mathrm{ICR}_i^{\lambda^k}(\tau_i)$ 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 et al. (2018) 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 et al. (2004), which unlike standard level-k reasoning, allows for uncertainty on the level of rationality attached to opponents. Kets (2014) and Heifetz and Kets (2018) 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.

[11] This 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,

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

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 $\lim_{k\to\infty} \lambda_k$ to his opponents' having unlimited depth of reasoning.

Most of the results of this paper (Propositions 1, 2, and 3, and Proposition 7) apply to every sequence λ, so in particular, also for the families of solution concepts considered above (in particular, Proposition 7 provides an epistemic foundation for all of them within a standard epistemic framework). In Proposition 6, we will focus on a particular class of perturbations:

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

$$\Lambda^0 = \left\{ \lambda \in \Lambda \;\middle|\; \begin{array}{ll} \text{(i)} & \lim\limits_{n\to\infty} \lambda_n = 0, \\ \text{(ii)} & \lambda_n \geq \lambda_{n+1} \text{ for any } n \in \mathbb{N} \end{array} \right\}.$$

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 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.

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, $\mathrm{ICR}_i^\lambda(t_i) = \mathrm{ICR}_i^\lambda(\tau_i(t_i))$.*

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 $\mathrm{ICR}_i^\lambda : \mathcal{T}_i \rightrightarrows A_i$ is an upper-hemicontinuous correspondence. This means that behavior which is excluded by ICR^λ at a certain type will still be excluded at nearby belief hierarchies. This is similar to results shown for ordinary ICR and ϵ-ICR in Dekel et al. (2007).

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

Remark 1. Notice that, for any $n \in \mathbb{N}$, any player i, and any sequence of probabilities λ, the correspondence $\mathrm{ICR}^{\lambda}_{i,n} : T_i \rightrightarrows A_i$ has closed domain and is closed-valued; thus, Proposition 2 and the closed graph theorem imply that $\mathrm{Graph}(\mathrm{ICR}^{\lambda}_{-i,n})$ is closed and, therefore, measurable, justifying Definition 2.

Proposition 1 turns out to be helpful not only in simplifying the definition of ICR^{λ}, but also in the proof of the next 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 (i.e., 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 $\langle \mathcal{G}, \mathcal{T} \rangle$ be a Bayesian game. Then, for any player i and any type t_i, we have that*

(i) $\mathrm{ICR}_i(t_i) = \mathrm{ICR}^{\bar{1}}_i(t_i)$.

(ii) The correspondence given by $\lambda \mapsto \mathrm{ICR}^{\lambda}_i(t_i)$ is upper-hemicontinuous everywhere and continuous at $\lambda = \bar{1}$.

The last result in this section illustrates the connection between λ-rationalizability and ε-rationalizability: for any $\varepsilon > 0$, there exists some strictly positive amount of suspicion of lack of common belief in rationality, represented by $\lambda = \bar{p}$ with $p < 1$, such that for every player and every belief hierarchy, every λ-rationalizable action is also ε-rationalizable.

Proposition 4 (λ-rationalizability and ε-rationalizability). *Let $\langle \mathcal{G}, \mathcal{T} \rangle$ be a Bayesian game. Then, for any $\varepsilon > 0$, any $n \geq 0$, any player i and any type t_i, we have that, for every $p \geq 1/(1+\varepsilon/(2M))$, $\mathrm{ICR}^{\bar{p}}_{i,n}(t_i) \subseteq \varepsilon\text{-}\mathrm{ICR}_{i,n}(t_i)$.*

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

[12] A 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.

are twofold. Proposition 5 proves the robustness of the WY-discontinuity for sequences λ with components λ_n sufficiently close to 1; even under perturbation in common belief in rationality, if higher-order belief in rationality remains above some threshold p, unique selection arguments à la Weinstein and Yildiz (2007) still work. However, Proposition 6 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 (2018), 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 (2018) work and this paper is examined in Section 4.2, where we also discuss the relevance of our results to global games.

4.1 *The WY-discontinuity and common belief in rationality*

First, we show that the WY-discontinuity persists under perturbations in common belief in rationality that keep higher-order belief in rationality above some high enough threshold. To formalize this insight, we need to recall first the following refinement of ICR due to Chen et al. (2014b).

Definition 3 (Robust selection, cf. Definition 4 by Chen et al. 2014b). Let $\langle \mathcal{G}, \mathcal{T} \rangle$ be a Bayesian game. Then, for any player i and any type t_i we say that action a_i can be *robustly selected* for type t_i if there exists some $\varepsilon > 0$ and some sequence $(\tau_i^n)_{n \in \mathbb{N}}$ approaching $\tau_i(t_i)$ such that ε-$\mathrm{ICR}_i(\tau_i^n) = \{a_i\}$ for any $n \in \mathbb{N}$. Let $\mathrm{RS}_i(t_i)$ denote the set of actions that can be robustly selected for type t_i.

Obviously, it is possible that a type does not admit a robust selection; however, it follows from Weinstein and Yildiz's (2007) structure theorem and Proposition 5 by Chen et al. (2014b) that if the richness assumption is satisfied then the set of belief hierarchies that admit a robust selection is generic. We can now state our first main result.

Proposition 5 (WY-discontinuity for persistently high λ). *Let $\langle \mathcal{G}, \mathcal{T} \rangle$ be a Bayesian game with finite type space. Then there exists some $p < 1$ such that for any player i, any type t_i, any λ with $\lambda \geq \bar{p}$, and any $a_i \in \mathrm{RS}_i(t_i)$, there exists some convergent sequence $(\tau_i^n)_{n \in \mathbb{N}}$ approaching $\tau_i(t_i)$ such that $\mathrm{ICR}_i^\lambda(\tau_i^n) = \{a_i\}$ for all $n \in \mathbb{N}$.*

Thus, at any type admitting the WY-discontinuity (i.e., with multiple robustly selected actions), the discontinuity persists when the ICR concept is replaced by $\mathrm{ICR}^{\bar{p}}$, or by ICR^λ with $\lambda \geq \bar{p}$. That is, even under this more permissive solution concept, representing bounded rationality, unique selection procedures work and any refinement sharper than robust selection will fail to be robust. Since, as shown by Chen et al. (2014b), under the richness condition every action which is ε-ICR for some $\varepsilon < 0$ can be robustly selected, the theorem applies to such actions. Thus, for large enough λ as in the proposition, any refinement that makes a selection among strict equilibria will fail to be robust under ICR^λ. Indeed, we have the following.

Corollary 1 (Robust selection of strictly rationalizable actions for persistently high λ). *Let $\langle \mathcal{G}, \mathcal{T} \rangle$ be a Bayesian game, with finite type space, which satisfies the richness condition. Then, for any type t_i and any $a_i \in \varepsilon$-ICR(t_i) for some $\varepsilon < 0$, there exists some $p < 1$ such that for any λ with $\lambda \geq \bar{p}$, there exists some convergent sequence $(\tau_i^n)_{n\in\mathbb{N}}$ approaching $\tau_i(t_i)$ such that $\mathrm{ICR}_i^\lambda(\tau_i^n) = \{a_i\}$ for all $n \in \mathbb{N}$.*

We note here that the converse of Proposition 5 fails. The following simple example, where no action can be robustly selected but the conclusion of Proposition 5 holds, is inspired by Section 4 in Chen et al. (2014b). Consider a case where a player is insensitive to others' actions and has two actions which are tied for best reply. Proposition 3 of Chen et al. (2014b) shows that no robust selection is possible in such a case. The conclusion of our Proposition 5, though, is satisfied for both actions. Since λ now has no impact, the choice of p is irrelevant, and one can simply use any sequence where payoffs are perturbed in a consistent direction to select one of the actions.[13]

Results such as Proposition 5 fail if we allow a different weakening of common belief in rationality, where belief in rationality becomes very low at high orders.

Proposition 6 (No WY-discontinuity for vanishing λ). *Let $\langle \mathcal{G}, \mathcal{T} \rangle$ be a Bayesian game with finite type space. Then, for any $\varepsilon < 0$, any player i, any type t_i, and any λ with $\lambda_n \to 0$, there exists a neighborhood U of $\tau_i(t_i)$ such that ε-$\mathrm{ICR}_i(t_i) \subseteq \mathrm{ICR}_i^\lambda(\tau_i)$, for any $\tau_i \in U$.*

If, for instance, each ICR action of type t_i is actually a strict best reply for some belief, then the ICR and ε-ICR sets are identical at a type t_i for some $\varepsilon < 0$. Proposition 6 then implies that ICR_i^λ is continuous at t_i for any λ with $\lambda_n \to 0$. Notice that Propositions 5 and 6 provide contrasting cases of the impact of higher-order belief in rationality on the WY-discontinuty. On the one hand, Proposition 5 states that the WY-discontinuity remains under perturbations of that maintain higher-order belief in rationality above a high-enough threshold. On the other hand, Proposition 6 tells us that the WY-discontinuity vanishes under perturbations of a different kind: 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 reestablished (even when such perturbations are considered in the sense of the product topology).[14] 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) ε-ICR actions (for $\varepsilon < 0$) remain in $\mathrm{ICR}_{i,k}$ 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.

[13]We thank an anonymous referee for suggesting this example.

[14]Notice 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 $\lambda_n \to 0$ referenced in Proposition 6.

A corollary to Proposition 6 states that the generic uniqueness result of Weinstein and Yildiz (closely related to WY-discontinuity) also fails whenever $\lambda_n \to 0$ and there is any type with multiple ε-rationalizable actions for some $\varepsilon < 0$. Under these conditions, there is an open set of types in which every element admits multiple λ-rationalizable actions.

COROLLARY 2 (Nonrobustness of generic uniqueness). *Let $\langle \mathcal{G}, \mathcal{T} \rangle$ be a Bayesian game. Then, for any player i for which there exist some $\varepsilon < 0$ and some type t_i such that $|\varepsilon\text{-ICR}_i(t_i)| > 1$, and for any λ with $\lambda_n \to 0$, the following set is not dense:*

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

Proposition 6 is related to previous work showing that for finite n, ICR_n satisfies a form of lower-hemicontinuity. See, for instance, Proposition 2 of Chen et al. (2010). In the case, mentioned earlier, that λ consists of finitely many 1's followed by 0's, ICR^λ is equivalent to ICR_n for finite n.

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 5, 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 6, unique selection will be impossible for games close enough to the original game, because all actions 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 6 shows that, under certain arbitrarily *small* perturbations in common belief in rationality, the WY-discontinuity vanishes; that is, continuity of behavior is restored, even under almost common belief in rationality. Going back to the terminology of Alaoui and Penta (2016) and Friedenberg et al. (2018), the theorem departs from the standard model in Weinstein and Yildiz (2007) by introducing perturbations in common belief assumptions regarding players' *rationality bounds*. Strzalecki (2014) and Heifetz and Kets (2018) study the impact on the WY-discontinuity of perturbations in common belief assumptions regarding players' *cognitive bound*. Specifically, Heifetz and Kets (2018) provides 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).

Our Proposition 6 sheds light on both rationality bounds and cognitive bounds. As discussed in Section 3.2, for given λ the cognitive bound is $\sup\{n \in \mathbb{N} | \lambda_n > 0\}$ and the rationality bound is $\sup\{n \in \mathbb{N} | \lambda_n = 1\}$. When λ satisfies $\lambda_n > 0$ for every n, ICR^λ represents common belief in infinite depth of reasoning (synonymously, a cognitive bound of ∞, also called unbounded cognition.) Thus, the failure of the WY-discontinuity in Proposition 6, and, in particular, robust multiplicity, are consistent with common belief in unbounded cognition, since the proposition requires only that λ be a sequence converging to 0. Of course, if λ converges to 0, there must exist some m such that $\lambda_n < 1$ for every $n \geq m$, meaning that every players' *rationality* bound is finite. Within the framework of ICR^λ, we see a clear distinction between unbounded cognition, which in many cases allows robust multiplicity, and the much stronger unbounded *rationality*, which eliminates robust multiplicity and restores the WY-discontinuity.

4.2.3 *Alternate definition of* ICR^λ An anonymous referee proposed an interesting alternate definition of ICR^λ, with the same general motivating ideas. We will call it $\mathrm{ICR2}^\lambda$ and discuss its merits here. To determine $\mathrm{ICR2}^\lambda_k$, one applies elimination steps using the first k values of λ, but *reversed* from our order. That is, one first runs a round with confidence λ_k in remaining actions being played, then with confidence λ_{k-1}, etc. down to λ_1, for k total rounds of elimination. To determine $\mathrm{ICR2}^\lambda_{k+1}$, then requires an entirely different $(k+1)$-step process, starting with λ_{k+1}. As before, the infinite intersection of the $\mathrm{ICR2}^\lambda_k$ determines $\mathrm{ICR2}^\lambda$. This may sound like a surprising way to apply the sequence λ_k, but it has a nice motivation. The final step, using λ_1, ensures that all players best-respond to a belief assigning at least λ_1 to rational actions of the other players. The penultimate step, using λ_2, means that other players's actions are based on assigning probability at least λ_2 to opposing actions, and so on. As pointed out by the referee, this process reflects actions that are consistent with a natural infinite sequence of statements. For Player 1, the first statement would be that Player 1 is rational and believes with probability at least λ_1 that Player 2 is rational. The second is that Player 1 is rational, and believes with probability at least λ_1 that: Player 2 is rational *and* believes Player 1 is rational with probability at least λ_2.

Our definition involves a much simpler and more intuitive modification to the elimination process. Players analyze the game by first reducing the available actions for each player to actions that are sometimes a best reply. Then they reduce to best replies to the remaining actions, but assuming only λ_1 confidence that the previous step has been applied, and so on. Our concept is motivated by the idea of players who apply the deductive process of successive elimination rules with only partial confidence. Both concepts also have an epistemic foundation, in terms of players' higher-order beliefs about rationality, as distinct from the elimination process. The foundation for $\mathrm{ICR2}^\lambda$ reduces to a single infinite statement, while ours requires an infinite series of independent statements, so this gives ICR2 a rival claim to simplicity.

In two significant special cases that we have mentioned, the case of constant λ and of λ consisting of finitely many 1's followed by zeroes (cases (i)–(iii) in Section 3.2), the two concepts coincide. Indeed, the two concepts are close enough that replacing ICR$^\lambda$ with ICR2$^\lambda$ has no effect on the two main results, Propositions 5 and 6. The constant case suffices to imply that Proposition 5 is unaffected. Proposition 6 is unaffected because the proof hinges on the fact that $\lambda_n \to 0$ makes the elimination process effectively finite, and this idea applies for either concept.

5. Epistemic foundation of λ-rationalizability

Finally, we formally analyze the epistemic foundation of 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 et al. (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 (1989), allowing heterogeneous weights on higher-order beliefs. Common λ-belief serves as the base of our epistemic characterization result in Proposition 7, 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 et al. (2011), the epistemic analysis is based on epistemic hierarchies and performed in state space $\Omega = \mathcal{E} \times A \times \Theta$, where $\mathcal{E} = \prod_{i\in I} \mathcal{E}_i$. For each player i, we denote $\Omega_i = \mathcal{E}_i \times A_i$, and for each state ω, we will consider the following projections: $\omega_i = \mathrm{Proj}_{\Omega_i}(\omega)$, $e_i(\omega) = \mathrm{Proj}_{\mathcal{E}_i}(\omega)$, $a_i(\omega) = \mathrm{Proj}_{A_i}(\omega)$ and $\theta(\omega) = \mathrm{Proj}_{\Theta}(\omega)$. 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 $R_i = \{\omega \in \Omega | a_i(\omega) \in BR_i(e_{i,1}(\omega))\}$. As usual let $R = \bigcap_{i\in I} R_i$ and $R_{-i} = \bigcap_{i\in I} R_i$. Note that all of 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 (1989). 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,

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

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

$$\begin{tikzcd}
\mathcal{E}_i \arrow[r, dashed, "q_i"] \arrow[d, leftrightarrow, "\psi_i"'] & \mathcal{T}_i \arrow[d, leftrightarrow, "\varphi_i"] \\
\Delta(\mathcal{E}_{-i} \times A_{-i} \times \Theta) \arrow[r, dashed, "\bar{q}_i"] & \Delta(\mathcal{T}_{-i} \times \Theta)
\end{tikzcd}$$

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 $[q_i = \tau_i] = \{\omega \in \Omega | q_i(e_i(\omega)) = \tau_i\}$ 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 7 (Epistemic foundation of ICR^λ). *Let $\langle \mathcal{G}, \mathcal{T} \rangle$ be a Bayesian game and λ, a sequence of probabilities. Then λ-rationalizability characterizes rationality and common λ-belief in rationality; that is, for any player i and any type t_i it holds that*

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

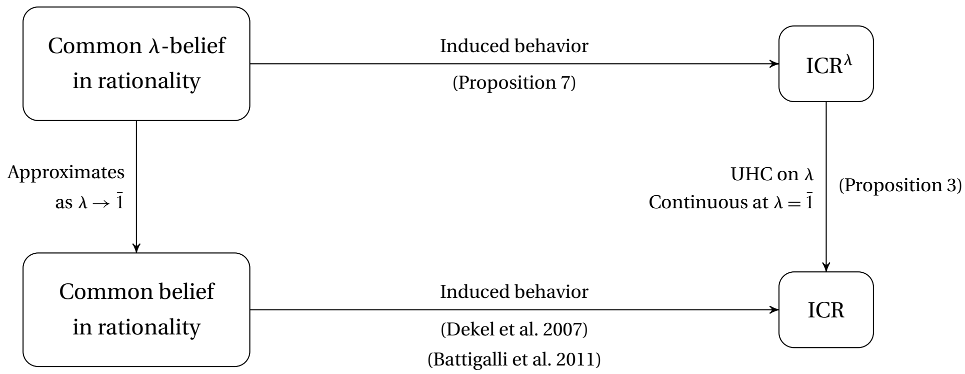

Figure 1. Rationalizability and perturbations in common belief in rationality.

The theoretical relevance of Proposition 7 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 nonconstant 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 not necessarily weighting different order belief in an homogeneous way (Proposition 7) and (ii) interim correlated λ-rationalizability is upper-hemicontinuous on λ and indeed, continuous when $\lambda = \bar{1}$ (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 et al. 2011, which corresponds to the $\lambda = \bar{1}$ 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.

Appendix A: Proofs: Properties of ICR^λ

A.1 *Elementary properties*

For convenience, we begin with the proof of Proposition 2.

Proposition 2 (Robustness to higher-order uncertainty about payoffs). *Let $\langle \mathcal{G}, \mathcal{T} \rangle$ be a Bayesian game. Then, for any $n \geq 0$, any player i, and any sequence of probabilities λ, correspondence $\mathrm{ICR}^\lambda_{i,n} : T_i \rightrightarrows A_i$ is upper-hemicontinuous. It follows that $\mathrm{ICR}^\lambda_i : T_i \rightrightarrows A_i$ is upper-hemicontinuous, too.*

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

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

for any $\ell = 1, \dots, n$ and, therefore, $\eta_i \in \mathrm{C}^\lambda_{i,n}(\tau_i)$. 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 $a_i \in \mathrm{ICR}^\lambda_{i,n+1}(\tau_i)$. □

An immediate corollary of this result is Remark 1, simply because closed sets are measurable. Remark 1 greatly simplifies the proof of the following lemma, providing an alternate characterization of ICR^λ, which is used in 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:*

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

Proof. We proceed by induction on n.

Initial step ($n = 1$). For the right-hand inclusion, pick $a_i \in \mathrm{ICR}^\lambda_{i,1}(\tau_i)$ and $\eta_i \in \mathrm{C}^\lambda_{i,0}(\tau_i)$ such that $a_i \in BR_i(\eta_i)$. Since $\mathrm{Proj}_{\mathcal{T}_{-i}\times\Theta} : \mathcal{T}_{-i} \times A_{-i} \times \Theta \to \mathcal{T}_{-i} \times \Theta$ is continuous and $\varphi_i(\tau_i)[E] = \eta_i[\mathrm{Proj}^{-1}_{\mathcal{T}_{-i}\times\Theta}(E)]$ for any measurable $E \subseteq \mathcal{T}_{-i} \times \Theta$, it follows immediately from the disintegration theorem that there exists a map $\sigma_{-i} : \mathcal{T}_{-i} \times \Theta \to \Delta(A_{-i})$ such that:[15]

(a) For each $E \subseteq A_{-i}$, map $\sigma^E_{-i} : \mathcal{T}_{-i} \times \Theta \to [0, 1]$ given $(\tau_{-i}, \theta) \mapsto \sigma_i(\tau_{-i}, \theta)[E]$ is measurable. Hence, σ_{-i} is measurable, too.[16]

[15]See Theorem 5.3.1 in Ambrosio et al. (2006, p. 121). We are working with compact and metrizable spaces; thus, in particular, all of them are Polish, and hence, Radon.

[16]Remember 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^E_{-i} 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.

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

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

(c)

$$\int_{A_{-i} \times \Theta} u_i\big((a_{-i}; a_i), \theta\big) \,\mathrm{d}(\mathrm{marg}_{A_{-i} \times \Theta} \eta_i)$$
$$= \int_{\mathcal{T}_{-i} \times \Theta} \bigg(\sum_{a_{-i} \in A_{-i}} \sigma_{-i}(\tau_{-i}, \theta)[a_{-i}] \cdot u_i\big((a_{-i}, a_i), \theta\big)\bigg) \,\mathrm{d}\varphi_i(\tau_i).$$

Then, since $\mathrm{ICR}^{\lambda}_{-i,0}(\tau_{-i}) = A_{-i}$ for and $\tau_{-i} \in \mathcal{T}_{-i}$, σ_{-i} obviously satisfies conditions (i) and (ii) in the statement of the lemma. For the left-hand inclusion, pick $a_i \in A_i$ and measurable $\sigma_{-i} : \mathcal{T}_{-i} \times \Theta \to \Delta(A_{-i})$ satisfying conditions (i) and (ii) above. Then define measure $\eta_i \in \Delta(\mathcal{T}_{-i} \times A_{-i} \times \Theta)$ as follows:

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

for any measurable $E \subseteq \mathcal{T}_{-i} \times A_{-i} \times \Theta$.[17] We now make the following two claims:

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

$$\eta_i[E \times A_{-i}] = \int_{\mathcal{T}_{-i} \times \Theta} \bigg(\sum_{a_{-i} \in A_{-i}} \sigma_{-i}(\tau_{-i}, \theta)\big[\mathrm{Proj}_{A_{-i}}\big(E \times A_{-i} \cap \big\{(\tau_{-i}, a_{-i}, \theta)\big\}\big)\big]\bigg) \,\mathrm{d}\varphi_i(\tau_i)$$
$$= \int_E \sigma_i(\tau_{-i}, \theta)[A_{-i}] \,\mathrm{d}\varphi_i(\tau_i) = \varphi_i(\tau_i)[E].$$

- $a_i \in BR_i(\eta_i)$. To see this, 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 $a'_i \in A_i$,

$$\int_{A_{-i} \times \Theta} u_i\big((a_{-i}; a'_i), \theta\big) \,\mathrm{d}(\mathrm{marg}_{A_{-i} \times \Theta} \eta_i)$$
$$= \sum_{a_{-i} \in A_{-i}} \int_{\mathcal{T}_{-i} \times \Theta} u_i\big((a_{-i}; a'_i), \theta\big) \,\mathrm{d}\nu_i(a_{-i})$$
$$= \int_{\mathcal{T}_{-i} \times \Theta} \bigg(\sum_{a_{-i} \in A_{-i}} \sigma_{-i}(\tau_{-i}, \theta)[a_{-i}] \cdot u_i\big((a_{-i}, a'_i), \theta\big)\bigg) \,\mathrm{d}\varphi_i(\tau_i).$$

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

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

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

Inductive step. Suppose that $n \geq 1$ is such that the claim holds. Let us check the $(n+1)$ case. For the right-hand inclusion, pick $a_i \in \mathrm{ICR}^{\lambda}_{i,n+1}(\tau_i)$ and $\eta_i \in \mathrm{C}^{\lambda}_{i,n}(\tau_i)$ such that $a_i \in BR_i(\eta_i)$, and family $(M_k)_{k=1}^n$ of measurable sets such that $M_k \subseteq \mathrm{Graph}(\mathrm{ICR}^{\lambda}_{-i,k})$ and $\eta_i[M_k] \geq \lambda_k$ for any $k = 1, \dots, n$. Then, since map $\mathrm{Proj}_{\mathcal{T}_{-i}\times\Theta} : \mathcal{T}_{-i} \times A_{-i} \times \Theta \to \mathcal{T}_{-i} \times \Theta$ is continuous and $\varphi_i(\tau_i)[E] = \eta_i[\mathrm{Proj}^{-1}_{\mathcal{T}_{-i}\times\Theta}(E)]$ for any measurable $E \subseteq \mathcal{T}_{-i} \times \Theta$, we know again from the disintegration theorem that there exists a map $\sigma_{-i} : \mathcal{T}_{-i} \times \Theta \to \Delta(A_{-i})$ 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 $k = 1, \dots, n$,

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

For the left-hand inclusion, pick $a_i \in A_i$ and measurable map $\sigma_{-i} : \mathcal{T}_{-i} \times \Theta \to \Delta(A_{-i})$ satisfying conditions (i) and (ii) for the $(n+1)$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:

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

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

- $\eta_i \in \mathrm{C}^{\lambda}_{i,0}(\tau_i)$. To see this, pick measurable $E \subseteq T_{-i} \times \Theta$ and develop

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

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

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

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

- $a_i \in BR_i(\eta_i)$. To see this, 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 $a'_i \in A_i$,

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

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

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

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

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,* $\mathrm{ICR}^{\lambda}_i(t_i) = \mathrm{ICR}^{\lambda}_i(\tau_i(t_i))$.

Proof. We will prove the slightly more general claim: for any player i, any type t_i, any sequence of probabilities λ, and any nonnegative integer n, it holds that $\mathrm{ICR}^{\lambda,\mathcal{T}}_{i,n}(t_i) = \mathrm{ICR}^{\lambda}_{i,n}(\tau_i(t_i))$. Let us 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 $k = 0, \ldots, n$, and fix $i \in I$, $t_i \in T_i$, and $\lambda \in \Lambda$. For the right inclusion, pick $a_i \in \mathrm{ICR}^{\lambda,\mathcal{T}}_{i,n+1}(t_i)$ and $\mu_i \in \mathrm{C}^{\lambda,\mathcal{T}}_{i,n}(t_i)$ such that $a_i \in BR_i(\mu_i)$, and for each $k = 1, \ldots, n$, $\mu_i[\mathrm{Graph}(\mathrm{ICR}^{\lambda,\mathcal{T}}_{-i,k}) \times \Theta] \geq \lambda_k$. Define now $\eta_i(\mu_i) \in \Delta(\mathcal{T}_{-i} \times A_{-i} \times \Theta)$ as follows:

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

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.[18] Notice that we have (i) that $\mathrm{marg}_{A_{-i}\times\Theta}\, \eta_i(\mu_i) = \mathrm{marg}_{A_{-i}\times\Theta}\, \mu_i$ and (ii) that[19]

$$\begin{aligned}
&\mathrm{marg}_{\mathcal{T}_{-i}\times\Theta}\, \eta_i(\mu_i)[E] = \eta_i(\mu_i)[A_{-i} \times E] \\
&\quad = \mu_i\big[\big\{(t_{-i}, a_{-i}, \theta) \in T_{-i} \times A_{-i} \times \Theta \big| \big(\big(\tau_{-i}(t_{-i}), a_{-i}, \theta\big) \in A_{-i} \times E\big\}\big] \\
&\quad = \mathrm{marg}_{T_{-i}\times\Theta}\, \mu_i\big[\big\{(t_{-i}, \theta) \in T_{-i} \times \Theta \big| \big(\tau_{-i}(t_{-i}), \theta\big) \in E\big\}\big] \\
&\quad = \pi_i(t_i)\big[\big\{(t_{-i}, \theta) \in T_{-i} \times \Theta \big| \big(\tau_{-i}(t_{-i}), \theta\big) \in E\big\}\big] \\
&\quad = \varphi_i\big(\tau_i(t_i)\big)[E].
\end{aligned}$$

[18]Due to every $\{(t_{-i}, a_{-i}, \theta) \in T_{-i} \times A_{-i} \times \Theta | ((\tau_{-i}(t_{-i})), a_{-i}, \theta) \in E\}$ being measurable.
[19]The fifth equality is a special case of formula (4) in Battigalli et al. (2011, p. 10).

Thus, it follows from (i) that $a_i \in BR_i(\eta_i(\mu_i))$, and from (ii) that $\eta_i(\mu_i) \in \mathrm{C}^{\lambda}_{i,0}(\tau_i(t_i))$. Now, fix $k = 0, \dots, n$ and note that we know, due to the induction hypothesis that[20]

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

Thus, we conclude that $\eta_i(\mu_i) \in \mathrm{C}^{\lambda}_{i,n}(\tau_i(t_i))$. For the left inclusion, we make use of Lemma 1. Pick $a_i \in \mathrm{ICR}^{\lambda}_{i,n+1}(\tau_i(t_i))$ and measurable $\sigma_{-i} : \mathcal{T}_{-i} \times \Theta \to \Delta(A_{-i})$ satisfying conditions (i) and (ii) in the statement of the lemma. Since map $f_{-i} : T_{-i} \times \Theta \to \mathcal{T}_{-i} \times \Theta$ 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:

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

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

- $\mu_i \in \mathrm{C}^{\lambda,\mathcal{T}}_{i,0}(t_i)$. To see this, pick measurable $E \subseteq T_{-i} \times \Theta$ and develop

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

- $\mu_i \in \mathrm{C}^{\lambda,\mathcal{T}}_{i,n}(t_i)$. Consider continuous map $F_i : T_{-i} \times A_{-i} \times \Theta \mapsto \mathcal{T}_{-i} \times A_{-i} \times \Theta$ given by $(t_{-i}, a_{-i}, \theta) \mapsto (f_{-i}(t_{-i}, \theta), a_{-i})$. Then we have that

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

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

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

$$\begin{aligned}&\int_{T_{-i} \times \Theta} \Bigg(\sum_{a_{-i} \in A_{-i}} \hat{\sigma}_{-i}(t_{-i}, \theta)[a_{-i}] \cdot u_i\big(\big(a_{-i}, a'_i\big), \theta\big)\Bigg)\, \mathrm{d}\pi_i(t_i) \\ &= \int_{\mathcal{T}_{-i} \times \Theta} \Bigg(\int_{\tau^{-1}_{-i}(\tau_{-i}) \times \{\theta\}} \Bigg(\sum_{a_{-i} \in A_{-i}} \hat{\sigma}_{-i}(t_{-i}, \theta)[a_{-i}] \cdot u_i\big(\big(a_{-i}, a'_i\big), \theta\big)\Bigg)\, \mathrm{d}\pi_i(t_i)\Bigg)\, \mathrm{d}\varphi_i(\tau_i) \\ &= \int_{\mathcal{T}_{-i} \times \Theta} \Bigg(\int_{\tau^{-1}_{-i}(\tau_{-i}) \times \{\theta\}} \Bigg(\sum_{a_{-i} \in A_{-i}} \sigma_{-i}\big(f_{-i}(t_{-i}), \theta\big)[a_{-i}] \cdot u_i\big(\big(a_{-i}, a'_i\big), \theta\big)\Bigg)\, \mathrm{d}\pi_i(t_i)\Bigg)\, \mathrm{d}\varphi_i(\tau_i)\end{aligned}$$

[20]Each $\mathrm{Graph}(\mathrm{ICR}^{\lambda,\mathcal{T}}_{-i,k})$ is clearly measurable; see Footnote 18.

$$= \int_{\mathcal{T}_{-i}\times\Theta} \left(\int_{\tau_{-i}^{-1}(\tau_{-i})\times\{\theta\}} \left(\sum_{a_{-i}\in A_{-i}} \sigma_{-i}(\tau_{-i},\theta)[a_{-i}] \cdot u_i\big((a_{-i},a_i'),\theta\big) \right) \mathrm{d}\pi_i(t_i) \right) \mathrm{d}\varphi_i(\tau_i)$$

$$= \int_{\mathcal{T}_{-i}\times\Theta} \left(\sum_{a_{-i}\in A_{-i}} \sigma_{-i}(\tau_{-i},\theta)[a_{-i}] \cdot u_i\big((a_{-i},a_i'),\theta\big) \right) \mathrm{d}\varphi_i(\tau_i).$$

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 $a_i' \in A_i$,

$$\int_{A_{-i}\times\Theta} u_i\big((a_{-i};a_i'),\theta\big)\,\mathrm{d}(\operatorname{marg}_{A_{-i}\times\Theta}\mu_i) =$$

$$= \sum_{a_{-i}\in A_{-i}} \int_{T_{-i}\times\Theta} u_i\big((a_{-i};a_i'),\theta\big)\,\mathrm{d}\nu_i(a_{-i})$$

$$= \int_{T_{-i}\times\Theta} \left(\sum_{a_{-i}\in A_{-i}} \hat{\sigma}_{-i}(t_{-i},\theta)[a_{-i}] \cdot u_i\big((a_{-i},a_i'),\theta\big) \right) \mathrm{d}\pi_i(t_i).$$

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

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

Proposition 3 (Robustness to higher-order uncertainty about rationality). *Let $\langle \mathcal{G}, \mathcal{T}\rangle$ be a Bayesian game. Then, for any player i and any type t_i, we have that*

(i) $\mathrm{ICR}_i(t_i) = \mathrm{ICR}^{\bar{1}}_i(t_i)$.

(ii) The correspondence given by $\lambda \mapsto \mathrm{ICR}^{\lambda}_i(t_i)$ is upper-hemicontinuous everywhere and continuous at $\lambda = \bar{1}$.

Proof. Since it follows immediately from Lemma 1 that for any $n \in \mathbb{N}$, any $i \in I$ and any $\tau_i \in \mathcal{T}_i$, $\mathrm{ICR}_{i,n}(\tau_i) = \mathrm{ICR}^{1}_{i,n}(\tau_i)$, 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 $\lambda \mapsto \mathrm{ICR}^{\lambda}_{i,n}(\tau_i)$ 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 $\mathrm{ICR}^{(\cdot)}_{-i,k}(\tau_{-i})$ is compact-valued, and hence, upper-hemicontinuity implies that $\bigcap_{n\geq 0} \mathrm{ICR}^{\lambda^n}_{-i,k}(\tau_{-i}) \subseteq \mathrm{ICR}^{\lambda}_{-i,k}(\tau_{-i})$ for any $(\lambda^n)_{n\in\mathbb{N}} \to \lambda$.[21] 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 $a_i^m \in \mathrm{ICR}^{\lambda^m}_{i,n+1}(\tau_i)$ for any $m \in \mathbb{N}$, and denote by (λ, a_i) the limit of the sequence. We need to check that $a_i \in \mathrm{ICR}^{\lambda}_{i,n+1}(\tau_i)$. First, pick $(\eta_i^m)_{\in\mathbb{N}} \in \prod_{m\in\mathbb{N}} \mathrm{C}^{\lambda^m}_{i,n}(\tau_i)$ 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

[21] Just write: $\Gamma(\lambda) = \mathrm{ICR}^{\lambda}_{-i,k}(\tau_{-i})$. 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, $\bigcap_{n\in\mathbb{N}} \Gamma(\lambda^n) \subseteq \Gamma(\lambda)$.

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 $\eta_i[\text{Graph}(\text{ICR}^{\lambda}_{-i,k})\times\Theta] \geq \lambda_k$ for any $k = 0,\dots,n$. Fix $k = 0,\dots,n$ and notice that $\eta_i^{m_\ell}[\text{Graph}(\text{ICR}^{\lambda^{m_\ell}}_{-i,k})\times\Theta] \geq \lambda_k^{m_\ell}$ 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 $A_{k,m_\ell} = \bigcup_{t\geq\ell}\text{Graph}(\text{ICR}^{\lambda^{m_t}}_{-i,k})$ 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) $\text{Graph}(\text{ICR}^{\hat{\lambda}^{m_t}}_{-i,k}) \subseteq \text{Graph}(\text{ICR}^{\hat{\lambda}^{m_\ell}}_{-i,k})$ for any $t \geq \ell$.

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

It follows from (i) and (ii) that $\overline{A_{k,m_\ell}} \subseteq \text{Graph}(\text{ICR}^{\hat{\lambda}^{m_\ell}}_{-i,k})$ for any $\ell \in \mathbb{N}$.[22] 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 $\eta_i[\text{Graph}(\text{ICR}^{\hat{\lambda}^{m_\ell}}_{-i,k})\times\Theta] \geq \hat{\lambda}_k^{m_\ell}$ 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

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

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

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

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

It follows from the above that for any $i \in I$, any $\tau_i \in \mathcal{T}_i$, and any $n \geq 0$, $\text{Graph}(\text{ICR}^{(\cdot)}_{i,n}(\tau_i))$ is closed, and thus, that so is $\text{ICR}^{\lambda}_{i,n}(\tau_i) = \text{Proj}_{A_i}((\{\lambda\}\times A_i)\cap\text{Graph}(\text{ICR}^{(\cdot)}_{i,n}(\tau_i)))$ for any $\lambda \in \Lambda$. Thus, $\text{ICR}^{(\cdot)}_{i,n}(\tau_i)$ is a compact-valued correspondence, and hence, by Theorem 17.25 in Aliprantis and Border (1999), we conclude that $\text{ICR}^{(\cdot)}_{i}(\tau_i) = \bigcap_{n\geq 0}\text{ICR}^{(\cdot)}_{i,n}(\tau_i)$ 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, Definition 17.2) that for any open subset $U \subseteq A_i$ such that $\text{ICR}^{\lambda}_{i}(\tau_i)\cap U \neq \varnothing$, there exists a neighborhood

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

$V \subseteq \Lambda$ of $\lambda = 1$ such that if $\lambda' \in V$, then $\mathrm{ICR}_i^{\lambda}(\tau_i) \cap U \neq \varnothing$. This follows immediately from the fact that, since $\lambda'_n \geq \lambda_n$ for any $n \in \mathbb{N}$, then $\mathrm{ICR}_i^{\lambda}(\tau_i) \subseteq \mathrm{ICR}_i^{\lambda'}(\tau_i)$. □

Proposition 4 (λ-rationalizability and ε-rationalizability). *Let $\langle \mathcal{G}, \mathcal{T} \rangle$ be a Bayesian game. Then, for any $\varepsilon > 0$, any $n \geq 0$, any player i and any type t_i, we have that, for every $p \geq 1/(1 + \varepsilon/(2M))$, $\mathrm{ICR}_{i,n}^{\bar{p}}(t_i) \subseteq \varepsilon\text{-}\mathrm{ICR}_{i,n}(t_i)$.*

Proof. Fix $\varepsilon > 0$, $i \in I$, $\tau_i \in \mathcal{T}_i$ and $p \geq 1/(1 + \varepsilon/(2M))$. We proceed by induction on n. The initial case ($n = 0$) holds trivially. Suppose that the claim is true for $n = k$; let us verify that then it also holds for $n = k + 1$. Fix $a_i \in \mathrm{ICR}_{i,k+1}^{\bar{p}}(\tau_i)$ and conjecture $\mu_i \in \mathrm{C}_{i,k}^{\bar{p}}(\tau_i)$ for which a_i is a best-reply. Now:

(i) If μ_i puts probability 1 on $\mathrm{Graph}(\mathrm{ICR}_{-i,k}) \times \Theta$, then set

$$\bar{\mu}_i = \mu_i.$$

(ii) If μ_i puts probability $q \in [p, 1)$ on $\mathrm{Graph}(\mathrm{ICR}_{-i,k}) \times \Theta$, then set

$$\bar{\mu}_i = \mu_i\big[\cdot \,|\, \mathrm{Graph}(\mathrm{ICR}_{-i,k}) \times \Theta\big] \quad \text{and}$$
$$\hat{\mu}_i = \mu_i\big[\cdot \,|\big(\mathrm{Graph}(\mathrm{ICR}_{-i,k}) \times \Theta\big)^c\big].$$

It follows from the induction hypothesis that in any case $\bar{\mu}_i \in \varepsilon\text{-}\mathrm{C}_{i,k}(\tau_i)$. Since $a_i \in BR_i(\mu_i)$ it follows that if (i) then, in particular, $a_i \in \varepsilon\text{-}BR_i(\mu_i)$, and thus, that $a_i \in \varepsilon\text{-}\mathrm{ICR}_{i,k+1}(\tau_i)$. If (ii) it follows from $a_i \in BR_i(\mu_i)$ that for any $a'_i \in A_i \setminus \{a_i\}$,

$$q \cdot \big(U_i(\bar{\mu}_i, a_i) - U_i(\bar{\mu}_i, a'_i)\big) + (1 - q) \cdot \big(U_i(\hat{\mu}_i, a_i) - U_i(\hat{\mu}_i, a'_i)\big) \geq 0,$$

and thus, that

$$\begin{aligned} U_i(\bar{\mu}_i, a_i) - U_i(\bar{\mu}_i, a'_i) &\geq -\left(\frac{1-q}{q}\right) \cdot \big(U_i(\hat{\mu}_i, a_i) - U_i(\hat{\mu}_i, a'_i)\big) \\ &\geq -\left(\frac{1-q}{q}\right) \cdot 2M \\ &\geq -\left(\frac{1-p}{p}\right) \cdot 2M \geq -\varepsilon. \end{aligned}$$

Then, since $a_i \in \varepsilon\text{-}BR_i(\bar{\mu}_i)$ and $\bar{\mu}_i \in \varepsilon\text{-}\mathrm{C}_{i,k}(\tau_i)$, we conclude that $a_i \in \varepsilon\text{-}\mathrm{ICR}_{i,k+1}(\tau_i)$ in this case, too. □

A.2 *Epistemic characterization*

Proposition 7 (Epistemic foundation of ICR^{λ}). *Let $\langle \mathcal{G}, \mathcal{T} \rangle$ be a Bayesian game and λ, a sequence of probabilities. Then λ-rationalizability characterizes rationality and common λ-belief in rationality; that is, for any player i and any type t_i it holds that*

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

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

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

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

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

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

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:

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

Since q_{-i} is a homeomorphism, η_i is well-defined, and obviously, it satisfies the following two conditions: (i) $\mathrm{marg}_{\mathcal{T}_{-i} \times \Theta}\, \eta_i = q_i(\tau_i)$ and (ii) $\mathrm{marg}_{A_{-i} \times \Theta}\, \eta_i = e_{i,1}(\omega)$. Thus, we have, first, that $\eta_i \in \mathrm{C}^{\lambda}_{i,0}(\tau_i)$, and, second, since $\omega \in R_i$, that $a_i \in BR_i(\eta_i)$. In consequence, $(\tau_i, a_i) \in \mathrm{Graph}(\mathrm{ICR}^{\lambda}_{i,1})$. 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 nonempty and has a 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 $a_i \in \mathrm{ICR}^{\lambda}_{i,1}(\tau_i)$, and $\eta_i \in \mathrm{C}^{\lambda}_{i,0}(\tau_i)$ 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:

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

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) $\mathrm{marg}_{\mathcal{T}_{-i} \times \Theta}\, \psi_i(e_i) = \mathrm{marg}_{\mathcal{T}_{-i} \times \Theta}\, \eta_i$ and (ii) $e_{i,1} = \mathrm{marg}_{A_{-i} \times \Theta}\, \eta_i$. Thus, it follows that $q_i(e_i) = \tau_i$ and $a_i \in BR_i(e_{i,1})$, and, therefore, that $(e_i, a_i) \in \mathrm{Proj}_{\mathcal{E}_i \times A_i}(R_i)$. Notice that, in particular, the proof of the right inclusion implies that $\Phi_{i,1}$ is nonempty.

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 nonempty 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:

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

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

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

Thus, $\eta_i \in C^{\lambda}_{i,k}(\tau_i)$ for any $k = 0,\dots,n+1$ and, in consequence, $(\tau_i,a_i)\in \mathrm{Graph}(\mathrm{ICR}^{\lambda}_{i,n+2})$. For the right inclusion, pick $(\tau_i,a_i)\in\mathcal{T}_i\times A_i$ such that $a_i\in\mathrm{ICR}^{\lambda}_{i,n+2}(\tau_i)$, and $\eta_i\in C^{\lambda}_{i,n+1}(\tau_i)$ such that $a_i \in BR_i(\eta_i)$. We know from the induction hypothesis that $\Phi_{j,n+1}$ is nonempty 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 us introduce the following notational convention: let $Z_{-i,k} = \mathrm{Graph}(\mathrm{ICR}^{\lambda}_{-i,k})$ and $W_{-i,k} = \mathrm{Proj}_{\Omega}(R_{-i}\cap B^{\lambda,k}_{-i}(R))$ 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:

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

where

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

for any $k = 0,\dots,n$. Since every $\phi_{-i,k+1}$ is measurable, $\psi_i(\eta_i)$ is well-defined. Set $e_i = \psi^{-1}_i(\psi_i(\eta_i))$ and $\omega_i = (e_i,a_i)$. Then we have that: (i) $\operatorname{marg}_{\mathcal{T}_{-i}\times\Theta}\psi_i(e_i) = \operatorname{marg}_{\mathcal{T}_{-i}\times\Theta}\eta_i$ and (ii) $e_{i,1} = \operatorname{marg}_{A_{-i}\times\Theta}\eta_i$. 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

$$\begin{aligned}
&\psi_i\big(e_i(\omega)\big)\big[\big\{\big(\omega'_{-i},\theta\big)\in\Omega_{-i}\times\Theta|\big(\omega'_{-i},\omega_i,\theta\big)\in R_{-i}\cap B^{\lambda,k}_{-i}(R)\big\}\big] \\
&\quad= \psi_i\big(e_i(\omega)\big)[\mathrm{Proj}_{\Omega_{-i}}W_{-i,k}\times\Theta] \\
&\quad= \sum_{\ell=0}^{n+1}\psi^{\ell}_i\big(e_i(\omega)\big)[\mathrm{Proj}_{\Omega_{-i}}W_{-i,k}\times\Theta]
\end{aligned}$$

$$
\begin{aligned}
&\geq \sum_{\ell=k+1}^{n} \eta_i\big[\big\{(\tau_{-i}, a_{-i}) \in (Z_{-i,\ell} \setminus Z_{-i,\ell+1}) | \big(\phi_{-i,\ell}(\tau_{-i}), a_{-i}\big) \in \mathrm{Proj}_{\Omega_{-i}} W_{-i,k}\big\} \times \Theta\big] \\
&\quad + \eta_i\big[\big\{(\tau_{-i}, a_{-i}) \in Z_{-i,n+1} | \big(\phi_{-i,n+1}(\tau_{-i}), a_{-i}\big) \in \mathrm{Proj}_{\Omega_{-i}} W_{-i,k}\big\} \times \Theta\big] \\
&\geq \sum_{\ell=k+1}^{n} \eta_i\big[(Z_{-i,\ell} \setminus Z_{-i,\ell+1}) \times \Theta\big] + \eta_i[Z_{-i,n+1} \times \Theta] \\
&= \eta_i[Z_{-i,k+1} \times \Theta] \geq \lambda_{k+1}.
\end{aligned}
$$

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 $(e_i, a_i) \in \mathrm{Proj}_{\mathcal{E}_i \times A_i}(R_i \cap B_i^{\lambda,n+1}(R))$ 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 nonempty.

Now, in order to complete the proof, fix $i \in I$ and $\tau_i \in \mathcal{T}_i$, and notice that for any $n \geq 0$,

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

Finally, the fact that

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

completes the proof. □

Appendix B: Proofs: Main results

B.1 *Robustness of the WY-discontinuity*

Proposition 5 (WY-discontinuity for persistently high λ). *Let $\langle \mathcal{G}, \mathcal{T} \rangle$ be a Bayesian game with finite type space. Then there exists some $p < 1$ such that for any player i, any type t_i, any λ with $\lambda \geq \bar{p}$, and any $a_i \in \mathrm{RS}_i(t_i)$, there exists some convergent sequence $(\tau_i^n)_{n \in \mathbb{N}}$ approaching $\tau_i(t_i)$ such that $\mathrm{ICR}_i^{\lambda}(\tau_i^n) = \{a_i\}$ for all $n \in \mathbb{N}$.*

Proof. Fix any t_i and use τ_i as an abbreviation for $\tau_i(t_i)$. Since a_i is robustly selected for t_i, we know that there exists some $\varepsilon > 0$ and some sequence $(\tau_i^n)_{n \in \mathbb{N}}$ converging to τ_i such that a_i is uniquely ε-rationalizable for τ_i^n for every $n \in \mathbb{N}$. Then it follows from Proposition 4 and nonemptiness of $\bar{p}$-rationalizability that for any $p \geq 1/(1 + \varepsilon/(2M))$,

a_i is uniquely $\bar{p}$-rationalizable for τ_i^n for every $n \in \mathbb{N}$. Then the result follows from the fact that ICR^λ is clearly monotone decreasing, i.e., for $\lambda \geq \bar{p}$ and every a_i, $\mathrm{ICR}_i^\lambda(a_i) \subseteq \mathrm{ICR}_i^{\bar{p}}(a_i)$. □

B.2 *Nonrobustness of the WY-discontinuity*

PROPOSITION 6 (No WY-discontinuity for vanishing λ). *Let $\langle \mathcal{G}, \mathcal{T} \rangle$ be a Bayesian game with finite type space. Then, for any $\varepsilon < 0$, any player i, any type t_i, and any λ with $\lambda_n \to 0$, there exists a neighborhood U of $\tau_i(t_i)$ such that ε-$\mathrm{ICR}_i(t_i) \subseteq \mathrm{ICR}_i^\lambda(\tau_i)$, for any $\tau_i \in U$.*

PROOF. Fix $i \in I$, let $t_i \in T_i$ and abbreviate $\tau_i(t_i)$ by simply τ_i. Let $a_i \in \varepsilon$-$\mathrm{ICR}_i(\tau_i)$ for some $\epsilon < 0$. We know from Proposition 2 in Chen et al. (2010) that for any $n \in \mathbb{N}$ there exists some $\delta_n > 0$ such that[23]

$$2\varepsilon\text{-}\mathrm{ICR}_{i,n}(\tau_i) \subseteq \varepsilon\text{-}\mathrm{ICR}_{i,n}(\hat{\tau}_i)$$

for any $\hat{\tau}_i \in U_i^n = B_{\delta_n}(\tau_i)$. Clearly, it follows that

$$a_i \in \varepsilon\text{-}\mathrm{ICR}_{i,n}(\hat{\tau}_i) \subseteq \mathrm{ICR}_{i,n}(\hat{\tau}_i)$$

for any $\hat{\tau}_i \in U_i^n$.

Now, since λ converges to 0 we know that there exists some $n_0 \in \mathbb{N}$ such that

$$\lambda_n \leq x := \frac{-\varepsilon}{2M - \varepsilon}$$

for any $n \geq n_0$. Then, for any $\hat{\tau}_i \in U_i$ take arbitrary $\hat{\eta}_i^1 \in \varepsilon$-$\mathrm{C}_{i,n_0}(\hat{\tau}_i)$ and arbitrary $\hat{\eta}_i^2 \in \bigcap_{n > n_0} \mathrm{C}_{i,n}(\hat{\tau}_i)$ and define:

$$\hat{\eta}_i = (1 - x) \cdot \hat{\eta}_i^1 + x \cdot \hat{\eta}_i^2.$$

Then:

- $a_i \in BR_i(\hat{\eta}_i)$. To see this, simply notice that for any $a_i' \in A_i \setminus \{a_i\}$ we have

$$\begin{aligned} U_i(\hat{\eta}_i, a_i) - U_i(\hat{\eta}_i, a_i') &= (1 - x) \cdot \big(U_i(\hat{\eta}_i^1, a_i) - U_i(\hat{\eta}_i^1, a_i')\big) \\ &\quad + x \cdot \big(U_i(\hat{\eta}_i^2, a_i) - U_i(\hat{\eta}_i^2, a_i')\big) \\ &\geq (1 - x) \cdot (-\varepsilon) + x \cdot 2M = 0. \end{aligned}$$

- $\hat{\eta}_i \in \bigcap_{n \geq 0} \mathrm{C}_{i,n}^\lambda(\hat{\tau}_i)$. Obviously, $\hat{\eta}_i \in \mathrm{C}_{i,0}^\lambda(\hat{\tau}_i)$. Now, for any natural $n \leq n_0$ we have

$$\hat{\eta}_i\big[\mathrm{Graph}\big(\mathrm{ICR}_{-i,n}^\lambda\big) \times \Theta\big] \geq \hat{\eta}_i\big[\mathrm{Graph}(\varepsilon\text{-}\mathrm{ICR}_{-i,n}) \times \Theta\big] = (1 - x) + x \geq \lambda_n.$$

  Whereas for $n > n_0$ we have

$$\hat{\eta}_i\big[\mathrm{Graph}\big(\mathrm{ICR}_{-i,n}^\lambda\big) \times \Theta\big] \geq x \cdot \hat{\eta}_i^2\big[\mathrm{Graph}(\mathrm{ICR}_{-i,n}) \times \Theta\big] = x \geq \lambda_n.$$

[23]Proposition 2 by Chen et al. (2010) is stated for $\varepsilon > 0$. It is routine to verify that the authors' proof generalizes to arbitrary ε. For better comparison between that result and this claim, suppose that $\gamma = 2\varepsilon$ and $\varepsilon = \gamma + 4M\delta_n$ and then take $\delta_n = -\varepsilon/4M$.

Thus, we conclude that $a_i \in \mathrm{ICR}_i^{\lambda}(\hat{\tau}_i)$ for every $\hat{\tau}_i \in U_i^{n_0}$. The fact that the action sets are finite ensures that U_i, the intersection of all the $U_i^{n_0}$ corresponding to each rationalizable action, is open too, and thus, the proof is complete. □

Corollary 2 (Nonrobustness of generic uniqueness). *Let $\langle \mathcal{G}, \mathcal{T} \rangle$ be a Bayesian game. Then, for any player i for which there exist some $\varepsilon < 0$ and some type t_i such that $|\varepsilon\text{-}\mathrm{ICR}_i(t_i)| > 1$, and for any λ with $\lambda_n \to 0$, the following set is not dense:*

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

Proof. This follows directly from Proposition 6: fix sequence λ with limit 0 and pick $i \in I$ and $t_i \in T_i$ such that $|\varepsilon\text{-}\mathrm{ICR}_i(t_i)| > 1$ for some $\epsilon < 0$. Then we know that there exists some open neighborhood U of t_i such that $\varepsilon\text{-}\mathrm{ICR}_i(t_i) \nsubseteq \mathrm{ICR}_i^{\lambda}(\hat{t}_i)$ for any $\hat{t}_i \in U$. Thus, the set $\mathcal{U}_i$ does not intersect open set U, and is not dense. □

References

Alaoui, Larbi and Antonio Penta (2016), "Endogenous depth of reasoning." *The Review of Economic Studies*, 83, 1297–1333. [98, 103]

Aliprantis, Charambolos D. and Kim C. Border (1999), *Infinite Dimensional Analysis: A Hitchhiker's Guide*, second edition. Springer-Verlag, Berlin. [114]

Ambrosio, Luigi, Nicola Gigli, and Giuseppe Savaré (2006), *Gradient Flows in Metric Spaces and in the Space of Probability Measures*. Lectures in Mathematics. ETH Zürich. Birkhäuser Verlag, Basel. [108]

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. [95, 96, 105, 106, 107, 111]

Brandenburger, Adam and Eddie Dekel (1993), "Hierarchies of beliefs and common knowledge." *Journal of Economic Theory*, 59, 189–198. [94, 95, 105]

Camerer, Colin F., George Loewenstein, and Matthew Rabin (2004), *Advances in Behavioral Economics*. The Roundtable Series in Behavioral Economics. Russel Sage Foundation, New York; Princeton University Press, New Jersey. [98]

Carlsson, Hans and Eric van Damme (1993), "Global games and equilibrium selection." *Econometrica*, 61, 989–1018. [93, 103]

Chen, Yi-Chun (2012), "A structure theorem for rationalizability in the normal form of dynamic games." *Games and Economic Behavior*, 75, 587–597. [93]

Chen, Yi-Chun, Alfredo Di Tillio, Eduardo Faingold, and Siyang Xiong (2010), "Uniform topologies on types." *Theoretical Economics*, 5, 445–478. [90, 93, 103, 119]

Chen, Yi-Chun, Alfredo Di Tillio, Eduardo Faingold, and Siyang Xiong (2017), "Characterizing the stratrategic impact of misspecified beliefs." *The Review of Economic Studies*, 84, 1424–1471. [93]

Chen, Yi-Chun, Satoru Takahashi, and Siyang Xiong (2014a), "The Weinstein–Yildiz critique and robust predictions with arbitrary payoff uncertainty." Unpublished paper, Department of Economics, National University of Singapore. [93]

Chen, Yi-Chun, Satoru Takahashi, and Siyang Xiong (2014b), "The robust selection of rationalizability." *Journal of Economic Theory*, 151, 448–475. [93, 101, 102]

Dekel, Eddie, Drew Fudenberg, and Stephen Morris (2006), "Topologies on types." *Theoretical Economics*, 1, 275–309. [93, 95, 96]

Dekel, Eddie, Drew Fudenberg, and Stephen Morris (2007), "Interim correlated rationalizability." *Theoretical Economics*, 2, 15–40. [95, 96, 98, 99]

Friedenberg, Amanda, Willemien Kets, and Terri Kneeland (2018), "Is bounded rationality driven by limited ability?" Unpublished paper, W. P. Carey School of Business, Arizona State University. [98, 103]

Germano, Fabrizio and Peio Zuazo-Garin (2017), "Bounded rationality and correlated equilibria." *International Journal of Game Theory*, 46, 595–629. [100]

Harsanyi, John C. (1967), "Games with incomplete information played by "Bayesian" players, part I, the basic model." *Management Science*, 14, 320–334. [93, 94]

Harsanyi, John C. (1968a), "Games with incomplete information played by "Bayesian" players, part II, Bayesian equilibrium points." *Management Science*, 14, 159–182. [93]

Harsanyi, John C. (1968b), "Games with incomplete information played by "Bayesian" players, part III, the basic probability distribution of the game." *Management Science*, 14, 486–502. [93]

Heifetz, Aviad and Dov Samet (1998), "Topology-free typology of beliefs." *Journal of Economic Theory*, 82, 324–341. [108]

Heifetz, Aviad and Willemien Kets (2018), "Robust multiplicity with a grain of naiveté." *Theoretical Economics*, 13, 415–465. [92, 93, 98, 101, 103]

Hu, Tai-Wei (2007), "On p-rationalizability and approximate common certainty of rationality." *Journal of Economic Theory*, 136, 379–391. [97, 105, 107]

Kajii, Atsushi and Stephen Morris (1997), "The robustness of equilibria to incomplete information." *Econometrica*, 65, 1283–1309. [93]

Kets, Willemien (2014), "Finite depth of reasoning and equilibrium play." Discussion Paper 1569, Center for Mathematical Studies in Economics and Management Science, Northwestern University. [98]

Monderer, Dov and Dov Samet (1989), "Approximating common knowledge with common beliefs." *Games and Economic Behavior*, 1, 170–190. [97, 105]

Monderer, Dov and Dov Samet (1996), "Proximity of information in games with incomplete information." *Mathematics of Operations Research*, 21, 707–725. [93]

Morris, Stephen and Hyun Song Shin (1998), “Unique equilibrium in a model of self-fulfilling currency attacks.” *American Economic Review*, 88, 587–597. [93]

Nagel, Rosemarie (1995), “Unraveling the guessing game: An experimental study.” *American Economic Review*, 85, 1313–1326. [98]

Penta, Antonio (2012), “Higher order uncertainty and information: Static and dynamic games.” *Econometrica*, 80, 631–660. [93]

Penta, Antonio (2013), “On the structure of rationalizability for arbitrary spaces of uncertainty.” *Theoretical Economics*, 8, 405–430. [93, 96]

Rubinstein, Ariel (1989), “The electronic mail game: Strategic behavior under ‘almost common knowledge’.” *American Economic Review*, 79, 385–391. [89, 91, 92]

Stahl, Dale O. and Paul W. Wilson (1994), “Experimental evidence on players’ models of other players.” *Journal of Economic Behavior & Organization*, 25, 309–327. [98]

Strzalecki, Tomasz (2014), “Depth of reasoning and higher order beliefs.” *Journal of Economic Behavior & Organization*, 108, 108–122. [92, 103]

Weinstein, Jonathan and Muhamet Yildiz (2007), “A structure theorem for rationalizability with application to robust predictions of refinements.” *Econometrica*, 75, 365–400. [89, 90, 92, 93, 96, 100, 101, 103]

Co-editor George J. Mailath handled this manuscript.

Manuscript received 16 December, 2016; final version accepted 22 June, 2019; available online 24 June, 2019.