Theoretical Economics 8 (2013), 405–430 1555-7561/20130405

On the structure of rationalizability for arbitrary spaces of uncertainty

Antonio Penta
Department of Economics, University of Wisconsin–Madison

Weinstein and Yildiz (2007) show that only very weak predictions are robust to misspecifications of higher order beliefs. Whenever a type has multiple rationalizable actions, any of these actions is uniquely rationalizable for some arbitrarily close type. Hence, refinements of rationalizability are not robust. This negative result is obtained under a *richness condition*, which essentially means that *all* common knowledge assumptions on payoffs are relaxed.

In many settings, this condition entails an unnecessarily demanding robustness test. It is, therefore, natural to explore the structure of rationalizability when arbitrary common knowledge assumptions are relaxed (i.e., without assuming *richness*).

For arbitrary spaces of uncertainty and for every player i, I construct a set $\mathcal{A}_i^\infty$ of actions that are uniquely rationalizable for some hierarchy of beliefs. The main result shows that for any type t_i and any action a_i rationalizable for t_i, if a_i belongs to $\mathcal{A}_i^\infty$ and is justified by conjectures concentrated on $\mathcal{A}_{-i}^\infty$, then there exists a sequence of types converging to t_i for which a_i is uniquely rationalizable. This result significantly generalizes Weinstein and Yildiz's. Some of its implications are discussed in the context of auctions and equilibrium refinements, and in connection with the literature on global games.

Keywords. Rationalizability, robustness, refinements, higher order beliefs, dominance solvability, richness, global games, structure theorems.

JEL classification. C72.

1. Introduction

Economic modelling naturally involves making common knowledge assumptions. Recently, Weinstein and Yildiz (2007) studied the robustness of game theoretic predictions when *all* such assumptions are relaxed. Assuming that the underlying space of uncertainty is sufficiently "rich," Weinstein and Yildiz prove a *structure theorem* for interim correlated rationalizability (ICR; Dekel et al. 2007) that has the following implications:

Antonio Penta: apenta@ssc.wisc.edu
This paper benefited from countless conversations with George Mailath. I am also grateful to Barton Lipman, Stephen Morris, Marion Oury, Dan Quint, Marzena Rostek, Bill Sandholm, Ricardo Serrano-Padial, Lones Smith, and the anonymous referees for helpful comments.

DOI: 10.3982/TE861

Result 1 (Nonrobustness). Whenever a model has multiple ICR outcomes, any of these outcomes is *uniquely* rationalizable in a model with beliefs arbitrarily close to the original ones.

Result 2 (Generic uniqueness). In the space of hierarchies of beliefs, the set of types with a unique ICR action is open and dense (i.e., models are generically *dominance-solvable*).

An important implication of Result 1 is that any refinement of ICR is not robust: any rationalizable outcome ruled out by some refinement of rationalizability would be uniquely selected (hence, *not* ruled out) in some arbitrarily close model of beliefs (that is, refinements of ICR are not upper hemicontinuous). Result 2 instead generalizes an important insight from the global games literature, that multiplicity is often the consequence of the common knowledge assumptions implicit in our models.

Weinstein and Yildiz's *richness condition* requires that for every action of every player, there exists a state under which that action is strictly dominant. This means that essentially *all* common knowledge assumptions are relaxed, as under the richness condition it is *not* common knowledge among the players that any action is *not* dominant. In many situations, this condition entails an unnecessarily demanding robustness test.

For instance, suppose that we are interested in the robust predictions of an auction model, $\mathcal{M}$. Following Weinstein and Yildiz, we would embed the original model (with common knowledge assumptions, $\mathcal{M}$) in a model $\mathcal{M}^*$ with a richer space of uncertainty, so as to be able to study the robustness question by looking at sequences of players' beliefs that converge to common belief in the original model $\mathcal{M}$. If model $\mathcal{M}^*$ satisfies the richness condition, Result 1 implies that the only robust predictions for $\mathcal{M}$ are those provided by ICR. But the richness condition requires the existence of dominance regions for all bids, including, for instance, those bids that (depending on the specific rules of the auction) are dominated for all types of bidders. But this essentially means that (in $\mathcal{M}^*$) it is no longer common knowledge that players are rational, that payoffs are decreasing in the price paid, or that players know the rules of the auction, etc. How would Result 1 change if $\mathcal{M}^*$, for instance, contained only dominance regions for bids consistent with the assumption of common knowledge of the auction's rules and of players' rationality?

More generally, in specific contexts, one may wish to investigate the robustness question when some common knowledge assumptions are relaxed, but not others. The reliance on the richness condition therefore significantly limits the applicability of the results above to specific applied models. It is, therefore, an important theoretical problem to investigate the robustness of solution concepts under arbitrary common knowledge restrictions. This paper explores the structure of ICR on arbitrary spaces of uncertainty, so as to analyze the robustness of game theoretic predictions to the relaxation of arbitrary common knowledge assumptions.

Fix a space of uncertainty Θ, let $\mathcal{A}_i^0$ be the set of actions of player i that are dominant in some state in Θ, and let $\mathcal{A}_{-i}^0 = \times_{j\neq i}\mathcal{A}_j^0$. Recursively, let $\mathcal{A}_i^k$ denote the set of actions of player i that are unique best responses to conjectures concentrated on $\Theta \times \mathcal{A}_{-i}^{k-1}$. Finally, define $\mathcal{A}_i^\infty = \bigcup_{k\geq 0}\mathcal{A}_i^k$. Theorem 1 (Section 3) shows the following result.

Result 1′. For any type (or hierarchy of beliefs) t_i and for every action a_i that is ICR for type t_i, if a_i belongs to $\mathcal{A}_i^\infty$ and is justified by conjectures concentrated on $\times_{j\neq i}\mathcal{A}_j^\infty$, then there exists a sequence of types converging to t_i along which a_i is uniquely rationalizable.

By requiring that every action a_i is strictly dominant in some state $\theta^{a_i}\in\Theta$, Weinstein and Yildiz's richness condition trivially implies that $\mathcal{A}_i^\infty = A_i$ for each i. In this case, Result 1 follows immediately from Result 1′. For arbitrary spaces of uncertainty, $\mathcal{A}_i^\infty$ generally is a subset of A_i, but in some cases, it may be that $\mathcal{A}_i^\infty = A_i$ even if richness is not assumed. In fact, it is shown that very mild relaxation of common knowledge assumptions suffice to guarantee that $\mathcal{A}^\infty = A$. Hence, the implications of Weinstein and Yildiz's results remain valid for significantly less demanding robustness tests, which reinforces their message.

The dominance states in Weinstein and Yildiz (2007) are often interpreted as "artificial" states that are added to the underlying space of uncertainty to model the relaxation of common knowledge assumptions. Under this interpretation, their exercise is similar to the introduction of "behavioral types" common in the literature on reputation and on strategic foundations of the rational expectations equilibrium.[1] Similarly, the results of this paper can be used to perform robustness exercises in which artificial dominance states are added for only *some* (as opposed to *all*) of the actions.

However, and more importantly, Theorem 1 can be applied to perform robustness exercises that do not involve the introduction of *any* artificial state. Clearly, in some cases, this enables stronger robust predictions than allowed by Weinstein and Yildiz's original result. However, results as powerful as Result 1 are possible even without introducing *any* artificial state. Consider the following example, which is more extensively discussed in Section 4.1 below.

Example (The "public good demand game"). A society must decide on the quantity of a public good, $x\in X=[0,1]$. The decision is made according to the following protocol: every agent proposes a quantity and the average of the proposals is implemented.

Agents' preferences depend on the realization of a state of nature, $\theta\in\Theta\subseteq\mathbb{R}$, represented by utility functions $U_i: X\times\Theta\to\mathbb{R}$ such that $U_i(x,\theta) = -(\theta - x)^2$. For simplicity, assume that $A_i = \Theta = \{0, \frac{1}{4}, \frac{1}{2}, \frac{3}{4}, 1\}$ and that the society is made of two individuals.

It can be checked (see Example 1, Section 4.1) that actions $a_i = 0, 1$ are the only ones that admit dominance states in Θ. We thus have $\mathcal{A}_i^0 = \{0,1\}$ and the richness condition is not satisfied. Nonetheless, it can be verified that $\mathcal{A}_i^\infty = A_i$. Hence, with no need to introduce *any* artificial state, Theorem 1 delivers the same nonrobustness result as Result 1 in this context. ♢

In addition to extending the robustness analysis to contexts in which the richness condition may not be compelling (see, e.g., Section 4.1), Theorem 1 provides a flexible

[1] In those literatures, such behavioral types are often referred to as *commitment*, *crazy*, *noise traders*, etc. (See references in footnote 16.)

tool that can be easily adapted to obtain further general results. Propositions 1, 2, and 3 in Section 4 represent some examples of such applications of Theorem 1.

Proposition 1 (Section 4.2) shows how an analogue of Result 1 can be obtained, when ICR is replaced by *arbitrary* equilibrium concepts: the result, which follows trivially from Theorem 1, holds under the assumption that the underlying space of uncertainty contains regions of dominance for the equilibrium actions only, which means that higher order beliefs do not rule out the possibility that the opponents may be "committed" to specific equilibrium actions (a very mild perturbation of common knowledge assumptions).

Proposition 2 (Section 4.3) provides mild conditions on Θ that guarantee that $\mathcal{A}^{\infty} = A$, so that the full results of Weinstein and Yildiz (including Result 2) are obtained. In addition to providing a generic uniqueness result without richness, Proposition 2 can be further specialized (Lemma 5) to shed some light on the connection between Weinstein and Yildiz (2007) and the literature on global games. A common message of the two approaches is that multiplicity often arises from the common knowledge assumptions of our models. Once such assumptions are relaxed (e.g., assuming richness), models are typically dominance-solvable.[2] Both approaches exploit an infection argument on players' hierarchies of beliefs reminiscent of the logic of the *e-mail game* (Rubinstein 1991). The formal settings though are quite different and not directly comparable. At a substantive level, two main differences can be identified.

(i) The global games literature assumes that players' actions are ordered and payoff functions are supermodular (i.e., it is common knowledge that the game exhibits strategic complementarities). In contrast, Weinstein and Yildiz make no assumptions on agents' payoffs or on the action spaces.

(ii) Weinstein and Yildiz require the existence of dominance states *for each* action of each player. The global games literature instead requires only dominance regions for the extreme actions.

Proposition 2 obtains a generic uniqueness result that combines the positive features of both approaches (that is, without assuming richness or imposing common knowledge of supermodularity). Loosely speaking, it is sufficient to assume that it is *not* common knowledge that the game *does not* exhibit strategic complementarities plus regions of dominance for only two actions of every player (and some other technicalities).

Proposition 3 (Section 4.4), also obtained from Theorem 1, shows that for any action that is uniquely rationalizable for some hierarchy of beliefs and that is justified by opponents' actions that are uniquely rationalizable for some type, a "nearby uniqueness" result analogous to Result 1′ holds. Hence, whenever that nearby uniqueness result fails for one of these actions, then that action is not uniquely rationalizable *anywhere* in the universal type space.

[2]See Morris and Shin (2003) for a thorough survey of the literature on global games.

2. Game theoretic framework

I consider static games with payoff uncertainty, i.e., tuples $G = \langle I, \Theta, (A_i, u_i)_{i\in I}\rangle$, where $I = \{1, \dots, n\}$ is the set of players; for each $i \in I$, A_i is the set of actions and $u_i : A \times \Theta \to \mathbb{R}$ is i's payoff function, where $A := \times_{i\in I} A_i$ and Θ is a parameter space. Assume that the sets I and A are finite, and that Θ is a compact metric space. Players' hierarchies of beliefs over Θ are defined as usual (see Mertens and Zamir 1985): for each $i \in N$, let $Z_i^1 = \Delta(\Theta)$ denote i's *first order beliefs*, and for $k \geq 1$, define recursively

$$Z_{-i}^k = \times_{j\neq i} Z_j^k$$

and

$$Z_i^{k+1} = \{(t_i^1, \dots, t_i^{k+1}) \in Z_i^k \times \Delta(\Theta \times Z_{-i}^k) : \operatorname{marg}_{\Theta\times Z_{-i}^{k-1}} t_i^{k+1} = t_i^k\}.^3$$

Agent i's first order beliefs are elements of $\Delta(\Theta)$; an element of $\Delta(\Theta \times Z_{-i}^{k-1})$ is a Θ-*based k-order belief* for every $k > 1$. The set of (collectively coherent) Θ-*hierarchies* is defined by

$$T_i^* = \{(t_i^1, t_i^2, \dots) \in \times_{k\geq 1}\Delta(\Theta \times Z_{-i}^{k-1}) : (t_i^1, \dots, t_i^k) \in Z_i^k\ \forall k \geq 1\}.$$

Players' Θ-hierarchies are represented by means of type spaces: a *type space* (or *Bayesian model*) is a tuple $\mathcal{T} = \langle \Theta', (T_i, \tau_i)_{i\in I}\rangle$ s.t. $\Theta' \subseteq \Theta$ and for each $i \in I$, T_i is the (compact) set of *types* of player i, and the continuous function $\tau_i : T_i \to \Delta(\Theta' \times T_{-i})$ assigns to each type of player i his beliefs about Θ and the opponents' types. A type space is *finite* if $|\Theta' \times T| < \infty$. The set $\Theta' \subseteq \Theta$ denotes the set of states that is common certainty among the types in T. This, in principle, can be a strict subset of the fundamental space Θ.

Each type in a type space induces a Θ-hierarchy: For each $t_i \in T_i$, the first order beliefs induced by $t_i \in T_i$ are obtained by the map $\hat{\pi}_i^1 : T_i \to \Delta(\Theta)$ that is defined as follows: for every measurable $E \subseteq \Theta$,

$$\hat{\pi}_i^1(t_i)[E] = \tau_i(t_i)\big[\{(\theta, t_{-i}) \in \Theta \times T_{-i} : \theta \in E\}\big].$$

For $k > 1$, the induced k-order beliefs are obtained by mappings $\hat{\pi}_i^k : T_i \to \Delta(\Theta \times Z_{-i}^{k-1})$, which are defined recursively as follows: for every measurable $E \subseteq \Theta \times Z_{-i}^{k-1}$,

$$\hat{\pi}_i^k(t_i)[E] = \tau_i(t_i)\big[\big\{(\theta, t_{-i}) \in \Theta \times T_{-i} : (\theta, \hat{\pi}_i^{k-1}(t_i)) \in E\big\}\big].$$

The map $\hat{\pi}_i^* : T_i \to T_i^*$, which is defined by

$$t_i \mapsto \hat{\pi}_i^*(t_i) = (\hat{\pi}_i^1(t_i), \hat{\pi}_i^2(t_i), \dots),$$

assigns to each type in a type space the corresponding Θ-*hierarchy of beliefs*. From Mertens and Zamir (1985), we know that when T^* is endowed with the product topology, there is a homeomorphism

$$\boldsymbol{\phi}_i : T_i^* \longrightarrow \Delta(\Theta \times T_{-i}^*)$$

[3]For any set X, $\Delta(X)$ denotes the set of probability measures over X, endowed with the topology of weak convergence.

that preserves beliefs of all orders: for all $t_i^* = (t_i^1, t_i^2, \dots) \in T_i^*$,

$$\operatorname{marg}_{\Theta \times Z_i^{k-1}} \boldsymbol{\phi}_i(t_i^*) = t_i^k \quad \forall k \geq 1.$$

Hence, the tuple $\mathcal{T}^* = \langle \Theta, (T_i^*, \tau_i^*)_{i \in N} \rangle$ with $\tau_i^* = \boldsymbol{\phi}_i$ is a type space. It is referred to as the (Θ*-based*) *universal type space*. Furthermore, any *nonredundant type space* (i.e., such that $\forall t_i, t_i' \in T_i$, $t_i \neq t_i'$ implies $\hat{\pi}_i^*(t_i) \neq \hat{\pi}_i^*(t_i')$) is a belief-closed subset of $\mathcal{T}^*$, in the sense that for every $\hat{\pi}_i^*(t_i) \in \hat{\pi}_i^*(T_i)$, we have $\boldsymbol{\phi}_i(\hat{\pi}_i^*(t_i))[\Theta \times \hat{\pi}_{-i}^*(T_{-i})] = 1$.[4] A *finite type* is any element $t_i \in T_i^*$ that belongs to a finite belief-closed subset of $\mathcal{T}^*$ (or, equivalently, any hierarchy that can be represented by means of a type in a finite type space). The set of finite types is denoted by $\hat{T}_i \subseteq T_i^*$.

Attaching a type space $\mathcal{T} = \langle \Theta', (T_i, \tau_i)_{i \in I} \rangle$ to the game with payoff uncertainty G, we obtain the Bayesian game $G^{\mathcal{T}} = \langle I, \Theta', (A_i, T_i, \hat{u}_i)_{i \in I} \rangle$, with payoff functions $\hat{u}_i : A \times \Theta' \times T \to \mathbb{R}$ s.t. $\hat{u}_i(a, \theta, t) = u_i(a, \theta)$ for all $(a, \theta, t) \in A \times \Theta' \times T$. Since players' types are payoff-irrelevant, with a slight abuse of notation, we write u_i and drop the dependence on T.[5]

Given a Bayesian game $G^{\mathcal{T}}$, player i's conjectures are denoted by $\psi^i \in \Delta(\Theta \times A_{-i} \times T_{-i})$.[6] For each type t_i, his *consistent conjectures* are

$$\Psi_i(t_i) = \{\psi^i \in \Delta(\Theta \times A_{-i} \times T_{-i}) : \operatorname{marg}_{\Theta \times T_{-i}} \psi^i = \tau_i(t_i)\}.$$

Let $\mathrm{BR}_i(\psi^i)$ denote the set of best responses to conjecture ψ^i:

$$\mathrm{BR}_i(\psi^i) = \arg\max_{a_i \in A_i} \sum_{(\theta, a_{-i}, t_{-i})} u_i(\theta, a_i, a_{-i}) \cdot \psi^i(\theta, a_{-i}, t_{-i}).$$

If $a_i \in \mathrm{BR}_i(\psi^i)$, we say that ψ^i *justifies* a_i. Appealing again to the payoff-irrelevance of the epistemic types, with another abuse of notation, we write $\mathrm{BR}_i(\psi^i)$ for conjectures $\psi^i \in \Delta(\Theta \times A_{-i})$ (i.e., disregarding the payoff-irrelevant component).

I present next the solution concept *interim correlated rationalizability* (ICR), introduced by Dekel et al. (2007).

Definition 1. Fix a Bayesian game $G^{\mathcal{T}}$. For each $i \in I$, let $\mathrm{ICR}_i^{\mathcal{T},0} = T_i \times A_i$. Recursively, for $k = 1, 2, \dots$ and $t_i \in T_i$, let $\mathrm{ICR}_{-i}^{\mathcal{T},k-1} = \times_{j \in I \setminus \{i\}} \mathrm{ICR}_j^{\mathcal{T},k-1}$,

$$\mathrm{ICR}_i^{\mathcal{T},k}(t_i) = \{a_i \in A_i : \exists \psi^{a_i} \in \Psi_i(t_i) \text{ s.t.: } a_i \in \mathrm{BR}_i(\psi^{a_i}) \text{ and}$$
$$\operatorname{supp}(\operatorname{marg}_{A_{-i} \times T_{-i}} \psi^{a_i}) \subseteq \mathrm{ICR}_{-i}^{\mathcal{T},k-1}\},$$

and $\mathrm{ICR}_i^{\mathcal{T},k} = \{(t_i, a_i) \in T_i \times A_i : a_i \in \mathrm{ICR}_i^{\mathcal{T},k}(t_i)\}$. Finally, $\mathrm{ICR}_i^{\mathcal{T}} := \bigcap_{k \geq 0} \mathrm{ICR}_i^{\mathcal{T},k}$ and $\mathrm{ICR}^{\mathcal{T}} = \times_{i \in I} \mathrm{ICR}_i^{\mathcal{T}}$.

[4]Given a function $f : X \to Y$, the set $f(X)$ is defined as $f(X) := \bigcup_{x \in X} f(x)$.

[5]In Weinstein and Yildiz (2007) and in the present setting, types are payoff-irrelevant or purely epistemic. This amounts to assuming that players have no information on payoffs (and this is common knowledge). See Penta (2012a) for the case with *payoff types* (i.e., with arbitrary information structures).

[6]Throughout the paper I maintain the convention that "beliefs" are about Θ and the opponents' beliefs about Θ; that is, beliefs are about exogenous variables only. The term "conjectures" instead refers to beliefs that also encompass the opponents' strategies.

Interim correlated rationalizability is a version of rationalizability (Pearce 1984 and Bernheim 1984) applied to the interim normal form, with the difference that the opponents' actions may be correlated (with one another and with the payoff state) in the eyes of a player.[7] Dekel et al. (2007) prove that whenever two types $t_i \in T_i$ and $t'_i \in T'_i$ are such that $\pi^*_i(t_i) = \pi^*_i(t'_i)$, then $\mathrm{ICR}^{\mathcal{T}'}_i(t'_i) = \mathrm{ICR}^{\mathcal{T}}_i(t_i)$; that is, ICR is completely determined by a type's hierarchies of beliefs, irrespective of the type space representation. Hence, we can drop the reference to the specific type space $\mathcal{T}$ and, without loss of generality, envision types as elements of the universal type space $\mathcal{T}^*$.

3. Structure of rationalizability without richness

Let $\mathcal{A}^0_i \subseteq A_i$ be the set of actions of player i for which there exists a *dominance state* $\theta^{a_i} \in \Theta$; that is, θ^{a_i} is such that a_i is strictly dominant.[8] Recursively, for each $k = 1, 2, \ldots,$ define

$$\mathcal{A}^k_i = \{a_i \in A_i : \exists \beta^i \in \Delta(\Theta \times \mathcal{A}^{k-1}_{-i}) \text{ s.t. } \{a_i\} = \mathrm{BR}_i(\beta^i)\}, \tag{1}$$

where $\mathcal{A}^k_{-i} = \times_{j\neq i}\mathcal{A}^k_j$ and $\mathcal{A}^k = \times_{i\in I}\mathcal{A}^k_i$, for each $k = 0, 1, \ldots$. Then let $\mathcal{A}^\infty_i = \bigcup_{k\geq 0}\mathcal{A}^k_i$.

In words, for each $k = 1, 2, \ldots,$ the set $\mathcal{A}^k_i$ consists of player i's actions that are a *unique* best responses to conjectures concentrated on $\mathcal{A}^{k-1}_{-i}$. Actions in $\mathcal{A}^0_i$ are those for which there exist dominance states. Then for each k, every action in $\mathcal{A}^k_i$ can be traced back to such dominance regions through a finite sequence of strict best responses.

Remark 1. It is easy to verify that for each $k = 1, 2, \ldots, \mathcal{A}^{k-1}_i \subseteq \mathcal{A}^k_i$. Also, since each A_i is finite, there exists $K \in \mathbb{N}$ such that for each $i \in I$, $\mathcal{A}^K_i = \mathcal{A}^{K+1}_i = \mathcal{A}^\infty_i$.

The main result is that for each t_i and for each action $a_i \in \mathrm{ICR}_i(t_i) \cap \mathcal{A}^\infty_i$ that is justified by conjectures concentrated on $\mathcal{A}^\infty_{-i}$, we can construct a sequence of (finite) types converging to t_i for which a_i is uniquely rationalizable.

Formally, let

$$\mathrm{ICR}_i(t_i; \mathcal{A}^\infty) = \{a_i \in \mathrm{ICR}_i(t_i) \cap \mathcal{A}^\infty_i : \exists \psi^{a_i} \in \Psi_i(t_i) \text{ s.t. } a_i \in \mathrm{BR}_i(\psi^{a_i}) \text{ and } \mathrm{supp}(\mathrm{marg}_{A_{-i}}\psi^{a_i}) \subseteq \mathcal{A}^\infty_{-i}\}. \tag{2}$$

Then the main result can be stated.

Theorem 1. *For each $t_i \in \hat{T}_i$ and for each $a_i \in \mathrm{ICR}_i(t_i; \mathcal{A}^\infty)$, there exists a sequence $\{t^\nu_i\} \subseteq \hat{T}_i$ s.t. $t^\nu_i \to t_i$ and for each $\nu \in \mathbb{N}$, $\{a_i\} = \mathrm{ICR}_i(t^\nu_i)$.*[9]

[7] Ely and Pęski (2006) study *interim (independent) rationalizability*, that is Pearce's solution concept applied to the interim normal form. Battigalli et al. (2011) study the relationships between these and other versions of rationalizability for incomplete information games.

[8] A slightly more general result can be obtained by letting $\mathcal{A}^0_i$ denote the set of player i's actions that are uniquely rationalizable for some state (see, e.g., Frankel et al. 2003). This point is discussed in Section 3.2. I maintain the dominance region terminology simply for expositional convenience.

[9] Since the universal type space is endowed with the product topology, this convergence as well as those in the following results are all with respect to that topology.

The proof of Theorem 1 is contained in Section 3.1. The uninterested reader may proceed directly to Sections 3.2 and 4, which contain examples and applications of the result.

3.1 *Proof of Theorem 1*

The next lemma shows that for each k and for each action $a_i \in \mathcal{A}_i^k$, there exists a finite type for which a_i is the only action that survives after $(k+1)$ rounds of iterated deletion of dominated actions.

Lemma 1. *For each $k = 0, 1, \ldots$, for each $a_i \in \mathcal{A}_i^k$ there exists a finite type $t_i' \in \hat{T}_i$ such that* $\mathrm{ICR}_i^{k+1}(t_i') = \{a_i\}$.

Proof. The proof is by induction.

Initial Step. This is immediate, as for $a_i \in \mathcal{A}_i^0$, there exists $\theta^{a_i} \in \Theta$ that makes a_i strictly dominant, and letting t_i' denote the type corresponding to common belief of θ^{a_i}, $\mathrm{ICR}_i^1(t_i') = \{a_i\}$.

Inductive Step. Let $a_i \in \mathcal{A}_i^k$. Then there exists $\beta^i \in \Delta(\Theta \times \mathcal{A}_{-i}^{k-1})$ such that $\{a_i\} = \mathrm{BR}_i(\beta^i)$. From the inductive hypothesis, there exists a function $\kappa_{-i}^{k-1} : \mathcal{A}_{-i}^{k-1} \to \hat{T}_{-i}$ such that for each $a_{-i} \in \mathcal{A}_{-i}^{k-1}$, $\{a_{-i}\} = \mathrm{ICR}_{-i}^k(\kappa_{-i}^{k-1}(a_{-i}))$. We want to show that there exists $t_i' \in \hat{T}_i$ such that $\mathrm{ICR}_i^{k+1}(t_i') = \{a_i\}$. Let $\mu^i \in \Delta(\Theta \times \mathcal{A}_{-i}^{k-1} \times \hat{T}_{-i})$ be defined by

$$\mu^i(\theta, a_{-i}, \kappa_{-i}^{k-1}(a_{-i})) = \beta^i(\theta, a_{-i})$$

and let t_i' be defined as $\tau_i^*(t_i') = \mathrm{marg}_{\Theta \times \hat{T}_{-i}} \mu^i$. Then, by construction,

$$\{\mu^i\} = \{\psi_i \in \Psi_i(t_i') : \mathrm{supp}(\mathrm{marg}_{A_{-i} \times \hat{T}_{-i}} \psi^i) \subseteq \mathrm{ICR}_{-i}^k\}$$

and

$$\{a_i\} = \mathrm{BR}_i(\mu^i).$$

Hence, $\mathrm{ICR}_i^{k+1}(t_i') = \{a_i\}$. □

Based on Lemma 1, the next definition introduces a set of types $\bar{T}_i \subseteq \hat{T}_i$ chosen so that each element of $\mathcal{A}_i^\infty$ is uniquely rationalizable for one type.

Definition 2. Let $\kappa_i^k : \mathcal{A}_i^k \to \hat{T}_i$ be defined as a mapping such that for each $a_i \in \mathcal{A}_i^k$, $\{a_i\} = \mathrm{ICR}_i^{k+1}(\kappa_i^k(a_i))$ and let $\kappa_i : \mathcal{A}_i^\infty \to \hat{T}_i$ be defined as a mapping such that for each $a_i \in \mathcal{A}_i^\infty$, $\{a_i\} = \mathrm{ICR}_i(\kappa_i(a_i))$. Given $\kappa_i : \mathcal{A}_i^\infty \to \hat{T}_i$, define the set of types $\bar{T}_i \subseteq \hat{T}_i$ as

$$\bar{T}_i := \{t_i \in \hat{T}_i : t_i = \kappa_i(a_i) \text{ for some } a_i \in \mathcal{A}_i^\infty\}.$$

(Notice that Definition 2 is well posed because of Lemma 1.)

Remark 2. Since $\mathcal{A}_i^\infty$ is finite, the set $\bar{T}_i$ is finite.

As already mentioned, Weinstein and Yildiz assume richness, that is, $\mathcal{A}_i^0 = A_i$ for each i. In that case, it is immediate to construct types with a unique rationalizable action. Given such dominance types, they prove their main result through an infection argument on players' hierarchies of beliefs. Their proof has two main steps: first, a type's beliefs are perturbed to show that any rationalizable action for that type is also strictly rationalizable for a nearby type (Lemma 6 in Weinstein and Yildiz 2007). Then, with a further perturbation, each strictly rationalizable action is made uniquely rationalizable for an arbitrarily close type, perturbing higher order beliefs only (Lemma 7 in Weinstein and Yildiz 2007).

With arbitrary spaces of uncertainty (i.e., without richness), the argument above requires two main modifications. First, the set of types $\bar{T}_i$ used to start the infection argument has to be constructed (Definition 2). Then a result analogous to Weinstein and Yildiz's Lemma 6 is proved (Lemma 3 below), but with a different solution concept than strict rationalizability, which is presented shortly (Definition 3).

The difference between Definition 3 and strict rationalizability parallels that between Weinstein and Yildiz's dominance types and the types in the set $\bar{T}_i$ constructed above. Given these preliminary steps, the further perturbations of higher order beliefs needed to obtain the result are completely analogous to Weinstein and Yildiz's: Lemma 4 below entails minor modifications of Weinstein and Yildiz's analogue (Lemma 7).

The proof of the main result is based on the following solution concept.

Definition 3. Given a type space $\langle \Theta, (T_i, \tau_i)_{i\in I}\rangle$, for each $i \in I$ and $t_i \in T_i$, let $\mathcal{W}_i^0 = T_i \times \mathcal{A}_i^0$. Recursively, for $k = 1, 2, \dots$, let $\mathcal{W}_{-i}^{k-1} = \times_{j\in I\setminus\{i\}} \mathcal{W}_j^{k-1}$, and for each $t_i \in T_i$,

$$\mathcal{W}_i^k(t_i) = \{a_i \in \mathcal{A}_i^k : \exists \psi^i \in \Delta(\Theta \times \mathcal{W}_{-i}^{k-1}) \text{ s.t. } \mathrm{marg}_{\Theta\times T_{-i}} \psi^i = \tau_i(t_i) \text{ and } \{a_i\} = \mathrm{BR}_i(\psi^i)\}$$

and $\mathcal{W}_i^k = \{(t_i, a_i) \in T_i \times A_i : a_i \in \mathcal{W}_i^k(t_i)\}$.

Let $K \in \mathbb{N}$ be such that for each $k \geq K$, $\mathcal{W}_i^{k+1}(t_i) \subseteq \mathcal{W}_i^k(t_i)$ for all t_i and i (such K exists because of Remark 1 above). Finally, define $\mathcal{W}_i(t_i) := \bigcap_{k\geq K} \mathcal{W}_i^k(t_i)$.

This solution concept is essentially an iterated deletion of never strict best responses for every type, except that for rounds $k < K$, the action sets are restricted to $\mathcal{A}_i^k$. Because the sets $\mathcal{A}_i^k$ are increasing for $k < K$, it is important to notice that $\mathcal{W}_i^k(t_i)$ may be nonmonotonic in k. Hence, up to K, $\mathcal{W}_i^k(t_i)$ may increase. When $k \geq K$ though, $\mathcal{A}_i^k = \mathcal{A}_i^\infty$ is constant and the condition $\exists \psi^i \in \Delta(\Theta \times \mathcal{W}_{-i}^{k-1})$ becomes (weakly) more and more stringent, making the sequence $\{\mathcal{W}_i^k(t_i)\}_{k>K}$ monotonically (weakly) decreasing. Being always nonempty, $\mathcal{W}_i^\infty(t_i) := \bigcap_{k>K} \mathcal{W}_i^k(t_i)$ is also nonempty and well defined (as long as $\mathcal{A}_i^0 \neq \varnothing$).

The following lemma states a standard fixed point property for $\mathcal{W}$; it is an immediate implication of Lemma 5 in Weinstein and Yildiz (2007).[10]

Lemma 2. *Given any type space $\langle \Theta, (T_i, \tau_i)_{i\in I}\rangle$, let $\{V_i(t_i)\}_{t_i\in T_i, i\in I}$ be a family of sets such that $V_i(t_i) \subseteq \mathcal{A}_i^\infty$ for all $i \in I$ and $t_i \in T_i$. If for each $a_i \in V_i(t_i)$, there exists*

[10]This is because $\mathcal{W}$ coincides with Weinstein and Yildiz's W^∞ applied to the game with actions $\mathcal{A}^\infty$.

$\psi^i \in \Delta(\Theta \times A_{-i} \times T^*_{-i})$ such that $\{a_i\} = \mathrm{BR}_i(\psi^i)$, $\mathrm{marg}_{\Theta\times T_{-i}}\psi^i = \tau_i(t_i)$, and $\psi^i(\theta, a_{-i}, t_{-i}) > 0 \Rightarrow a_{-i} \in V_{-i}(t_{-i})$, then $V_i(t_i) \subseteq \mathcal{W}_i(t_i)$ for each t_i.

We turn next to the analysis of higher order beliefs. The next lemma shows how, for each t_i and each action $a_i \in \mathrm{ICR}_i(t_i; \mathcal{A}^\infty)$, we can construct a sequence of (finite) types converging to t_i for which a_i survives the iterated deletion procedure introduced in Definition 3.

Lemma 3. *For each $a_i \in \mathrm{ICR}_i(t_i; \mathcal{A}^\infty)$, there exists $t_i(\varepsilon) \to t_i$ as $\varepsilon \to 0$ such that for each $\varepsilon > 0$, $a_i \in \mathcal{W}_i(t_i(\varepsilon))$ and $t_i(\varepsilon) \in \hat{T}_i$ (hence $a_i \in \mathcal{W}_i^k(t_i(\varepsilon))$ for all $k \geq K$).*

Proof. For each i, define the set $\mathrm{ICR}_i^{\mathcal{A}} = \{(t_i, a_i) : a_i \in \mathrm{ICR}_i(t_i; \mathcal{A}^\infty)\}$. For each $(t_i, a_i) \in \mathrm{ICR}_i^{\mathcal{A}}$, (2) implies that there exists $\psi^{a_i} \in \Delta(\Theta \times \mathrm{ICR}_{-i})$ such that $a_i \in \mathrm{BR}_i(\psi^i)$, $\mathrm{marg}_{\Theta\times T_{-i}}\psi^i = \tau_i(t_i)$, and $\mathrm{supp}(\mathrm{marg}_{A_{-i}}\psi^{a_i}) \subseteq \mathcal{A}^\infty_{-i}$. Furthermore, since $a_i \in \mathcal{A}^\infty_i$, there exists $\beta^i \in \Delta(\Theta \times \mathcal{A}^\infty_{-i})$ such that $\{a_i\} = \mathrm{BR}_i(\beta^i)$. Let κ_{-i} be as in Definition 2 and let $v_i^{(t_i,a_i)} \in \Delta(\Theta \times \bar{T}_{-i})$ be such that for each $(\theta, a_{-i}) \in \Theta \times \mathcal{A}^\infty_{-i}$, $v_i^{(t_i,a_i)}(\theta, \kappa_{-i}(a_{-i})) = \beta^i(\theta, a_{-i})$. For each $\varepsilon \in [0, 1]$, consider the type space $(T_i^\varepsilon, \tau_i^\varepsilon)_{i\in I}$ such that the set of types T_i^ε is given by $T_i^\varepsilon = \bar{T}_i \cup T_i^\iota$, where $\bar{T}_i$ is as in Definition 2 and T_i^ι is a finite set of types such that $T_i^\iota := \{\bar{\iota}_i^\varepsilon(t_i, a_i) : (t_i, a_i) \in \mathrm{ICR}_i^{\mathcal{A}}\}$; the beliefs $\tau_i^\varepsilon(\bar{t}_i^{a_i}) \in \Delta(\Theta \times \bar{T}_{-i})$ of types $\bar{t}_i^{a_i} \in \bar{T}_i$ are as in Definition 2 (i.e., such that $\mathrm{ICR}_i(\bar{t}_i^{a_i}) = \{a_i\}$), while the beliefs $\tau_i^\varepsilon(\bar{\iota}_i^\varepsilon(t_i, a_i)) \in \Delta(\Theta \times T_{-i}^\varepsilon)$ of types $\bar{\iota}_i^\varepsilon(t_i, a_i) \in T_i^\iota$ are defined as

$$\tau_i^\varepsilon(\bar{\iota}_i^\varepsilon(t_i, a_i)) = \varepsilon \cdot v_i^{(t_i,a_i)} + (1-\varepsilon)[\psi^{a_i} \circ \hat{\iota}_{-i,\varepsilon}^{-1}], \tag{3}$$

where $\hat{\iota}_{-i,\varepsilon} : \Theta \times \mathrm{ICR}_{-i}^{\mathcal{A}} \to \Theta \times T_{-i}^\varepsilon$ is the mapping given by $\hat{\iota}_{-i,\varepsilon}(\theta, a_{-i}, t_{-i}) = (\theta, \bar{\iota}_{-i}^\varepsilon(t_{-i}, a_{-i}))$ and $\psi^{a_i} \circ \hat{\iota}_{-i,\varepsilon}^{-1}$ denotes the measure on $\Theta \times T_{-i}^\varepsilon$ such that, for every measurable $E \subseteq \Theta \times T_{-i}^\varepsilon$,

$$(\psi^{a_i} \circ \hat{\iota}_{-i,\varepsilon}^{-1})(E) = \psi^{a_i}\big(\{(\theta, a_{-i}, t_{-i}) : \hat{\iota}_{-i,\varepsilon}(\theta, a_{-i}, t_{-i}) \in E\}\big).$$

Define $\gamma : \Theta \times T_{-i}^\varepsilon \to \Theta \times \mathcal{A}^\infty_{-i} \times T_{-i}^\varepsilon$ such that

$$\gamma(\theta, t_{-i}) = \begin{cases} (\theta, a_{-i}, \bar{\iota}_{-i}^\varepsilon(t_{-i}, a_{-i})) & \text{if } t_{-i} = \bar{\iota}_{-i}^\varepsilon(t_{-i}, a_{-i}) \in T_{-i}^\iota \\ (\theta, a_{-i}, \bar{t}_{-i}^{a_{-i}}) & \text{if } \bar{t}_{-i}^{a_{-i}} \in \bar{T}_{-i}. \end{cases}$$

Consider the conjectures $\psi^i \in \Delta(\Theta \times \mathcal{A}^\infty_{-i} \times T_{-i}^\varepsilon)$ defined by $\psi^i = (\tau_i^\varepsilon(\bar{\iota}_i^\varepsilon(t_i, a_i)) \circ \gamma^{-1})$. By construction, they are consistent with type $\bar{\iota}_i^\varepsilon(t_i, a_i)$. Furthermore, since ψ^i is a mixture of ψ^{a_i} (which makes a_i a best reply) and of β^i (which makes a_i a strict best reply), we have that $\{a_i\} = \mathrm{BR}_i(\psi^i)$. Hence, setting $V_i(\bar{\iota}_i^\varepsilon(t_i, a_i)) = \{a_i\}$ and $V_i(\bar{t}_i^{a_i}) = \{a_i\}$ in Lemma 2, we obtain $a_i \in \mathcal{W}_i(t_i)$ for $t_i = \bar{\iota}_i^\varepsilon(t_i, a_i), \bar{t}_i^{a_i}$. Finally, from (3) it is immediate to verify that $\bar{\iota}_i^\varepsilon(t_i, a_i) \to t_i$ as $\varepsilon \to 0$. □

The next lemma shows that for any type t_i and for any $a_i \in \mathcal{W}_i^k(t_i)$, $k = 0, 1, \ldots$, there exists a type that differs from t_i only for beliefs of order higher than k, for which a_i is the unique action that survives $(k+1)$ rounds of the ICR procedure.

For any type $t_i \in T_i^*$, let t_i^m denote the *mth order beliefs* of type t_i. (By definition of T_i^*, any $t_i \in T_i^*$ can be written as $t_i = (t_i^m)_{m=1}^{\infty}$.)

LEMMA 4. *For each $k = 0, 1, \ldots$ and for each $a_i \in \mathcal{W}_i^k(t_i)$, there exists $\tilde{t}_i = (\tilde{t}_i^m)_{m\in\mathbb{N}}$ such that $\tilde{t}_i^m = t_i^m$ for all $m \leq k$ and such that $\{a_i\} = \mathrm{ICR}_i^{k+1}(\tilde{t}_i)$.*

PROOF. The proof is by induction. For $k = 0$, $a_i \in \mathcal{W}_i^0(t_i) = \mathcal{A}_i^0$, so there exists, for action a_i, a dominance state θ^{a_i}. Let $\tilde{t}_i$ denote common belief of θ^{a_i}, so that $\{a_i\} = \mathrm{ICR}_i^1(\tilde{t}_i)$ (condition $\tilde{t}_i^0 = t_i^0$ holds vacuously). For the inductive step, write each t_{-i} as $t_{-i} = (l, h)$, where

$$l = (t_{-i}^1, \ldots, t_{-i}^k) \quad \text{and} \quad h = (t_{-i}^{k+1}, t_{-i}^{k+2}, \ldots).$$

Let

$$L = \{l : \exists h \text{ s.t. } (l, h) \in T_{-i}^*\}.$$

Let $a_i \in \mathcal{W}_i^k(t_i)$ and let $\psi^{a_i} \in \Delta(\Theta \times \mathcal{W}_{-i}^{k-1})$ be the corresponding conjecture s.t. $\mathrm{marg}_{\Theta\times T_{-i}}\psi^{a_i} = \tau_i(t_i)$ and $\{a_i\} = \mathrm{BR}_i(\psi^{a_i})$. Under the inductive hypothesis, for each $(a_{-i}, t_{-i}) \in \mathrm{supp}(\mathrm{marg}_{A_{-i}\times T_{-i}}\psi^{a_i})$, $\exists \tilde{t}_{-i}(a_{-i}) = (l, \tilde{h}(a_{-i}))$ s.t. $\mathrm{ICR}_{-i}^k(\tilde{t}_{-i}(a_{-i})) = \{a_{-i}\}$. Define the mapping

$$\varphi : \mathrm{supp}(\mathrm{marg}_{\Theta\times A_{-i}\times L}\psi^{a_i}) \to \Theta \times T_{-i}^*$$

by $\varphi(\theta, a_{-i}, l) = (\theta, \tilde{t}_{-i}(a_{-i}))$. Define $\tilde{t}_i$ by

$$\tau_i^*(\tilde{t}_i) = (\mathrm{marg}_{\Theta\times A_{-i}\times L}\psi^{a_i}) \circ \varphi^{-1}.$$

By construction,

$$\begin{aligned}
\mathrm{marg}_{\Theta\times A_{-i}\times L}\tau_i^*(\tilde{t}_i) &= \psi^{a_i} \circ \mathrm{proj}_{\Theta\times A_{-i}\times L}^{-1} \circ \varphi^{-1} \circ \mathrm{proj}_{\Theta\times L}^{-1} \\
&= \psi^{a_i} \circ \mathrm{proj}_{\Theta\times L}^{-1} \\
&= \psi^{a_i} \circ \mathrm{proj}_{\Theta\times A_{-i}\times T_{-i}^*}^{-1} \circ \mathrm{proj}_{\Theta\times L}^{-1} \\
&= \mathrm{marg}_{\Theta\times A_{-i}\times L}\tau_i(t_i),
\end{aligned}$$

where the first equality exploits the definition of lower order beliefs and the construction of type $\tilde{t}_i$, and the second equality follows from the definition of φ, for which

$$\mathrm{proj}_{\Theta\times L\times A_{-i}}^{-1} \circ \varphi^{-1} \circ \mathrm{proj}_{\Theta\times L}^{-1} = \mathrm{proj}_{\Theta\times L}^{-1}.$$

The third equality is simply notational and the last equality is by definition. Hence, by construction, we have $\mathrm{ICR}_i^{k+1}(\tilde{t}_i) = \{a_i\}$, which completes the inductive step. □

We are now in the position to complete the proof of Theorem 1, which we restate here.

THEOREM 1. *For each $t_i \in \hat{T}_i$ and for each $a_i \in \mathrm{ICR}_i(t_i; \mathcal{A})$, there exists a sequence $\{t_i^\nu\} \subseteq \hat{T}_i$ s.t. $t_i^\nu \to t_i$ and for each $\nu \in \mathbb{N}$, $\{a_i\} = \mathrm{ICR}_i(t_i^\nu)$.*

Proof. Take any $t_i \in \hat{T}$ and any $a_i \in \mathrm{ICR}_i(t_i, \mathcal{A}^\infty)$. From Lemma 3, there exists a sequence of finite types $t_i(\varepsilon) \to t_i$ (as $\varepsilon \to 0$) such that $a_i \in \mathcal{W}_i(t_i(\varepsilon))$ for each $\varepsilon > 0$, hence, there exists a sequence $\{t_i(n)\}_{n\in\mathbb{N}}$ converging to t_i such that $a_i \in \mathcal{W}_i^k(t_i(n))$ for all $k \geq K$. Then we can apply Lemma 4 to the types $t(n)$: for each n, for each $k \geq K$, and for each $a_i \in \mathcal{W}^k(t_i(n))$, there exists $\tilde{t}_i(k,n)$ such that $\tilde{t}_i^k(k,n) = t_i^k(n)$ and $\{a_i\} = \mathrm{ICR}_i^{k+1}(\tilde{t}_i(k,n))$. Hence, for each n, the sequence $\{\tilde{t}_i(k,n)\}_{k\in\mathbb{N}}$ converges to $t_i(n)$ as $k \to \infty$. Because the universal type space T^* is metrizable, there exists a sequence $t_i(n,k_n) \to t_i$ such that $\mathrm{ICR}_i(t_i(n,k_n)) = \{a_i\}$. Set $t_i^\nu = t_i(n,k_n)$: $t_i^\nu \to t_i$ as $\nu \to \infty$ and set $\mathrm{ICR}_i(t_i^\nu) = \{a_i\}$ for each ν. □

3.2 *Some comments on the main result*

For ease of reference, I list here some comments on Theorem 1. I further discuss these comments in the next section, which explores some of their implications.

1. Weinstein and Yildiz's richness condition amounts to assuming that Θ is such that $\mathcal{A}_i^0 = A_i$ for each i. In this case, Theorem 1 coincides with Proposition 1 in Weinstein and Yildiz (2007).

2. All results in Weinstein and Yildiz (2007) (including the generic uniqueness result) hold true, without richness, whenever $\mathcal{A}^\infty = A$. (One such possibility is explored in Section 4.3.)

3. Consider a Bayesian game $G^{\mathcal{T}}$. For every i, let $\mathcal{R}_i = \bigcup_{t_i \in T_i} \mathrm{ICR}(t_i)$ denote the set of rationalizable actions for player i. It follows immediately from Theorem 1 that if $\mathcal{A}_i^0 = \mathcal{R}_i$ for every i, then $\forall i$, $\forall t_i$, $\forall a_i \in \mathrm{ICR}_i(t_i)$, there exists a sequence $\{t_i^\nu\} \subseteq \hat{T}_i$ s.t. $t_i^\nu \to t_i$ and for each $\nu \in \mathbb{N}$, $\{a_i\} = \mathrm{ICR}_i(t_i^\nu)$.

4. Following Frankel et al. (2003), it is easy to see that Theorem 1 remains true (hence, so do all the subsequent comments) if $\mathcal{A}_i^0$ is defined as the set of actions for which there exist payoff states that make these actions uniquely rationalizable.[11]

4. Some extensions and applications

A Bayesian game is a commonly known tuple $G^{\mathcal{T}} = \langle N, \Theta', (A_i, T_i', \tau_i', u_i)_{i\in N}\rangle$. If the set Θ' satisfies the richness condition and if $\mathcal{T} = (T_i', \tau_i')_{i\in N}$ coincides with the Θ'-based universal type space, then maintaining common knowledge of $G^{\mathcal{T}}$ entails essentially no loss of generality. However, game theoretic modelling typically involves smaller (i.e., nonuniversal) type spaces, thereby imposing common knowledge assumptions on players' beliefs. Furthermore, specific applications often deliver common knowledge assumptions in the sense that the natural space of uncertainty Θ' does not satisfy the richness condition. A natural example is provided by situations in which players' payoffs $u_i : A \to \mathbb{R}$ are derived from players' preferences over some underlying space of outcomes, X, plus some rules of the game that specify agents' possible moves and the mapping that assigns outcomes to such moves (see Section 4.1).

[11] I am grateful to Stephen Morris for this remark.

Whenever the natural Θ' entails common knowledge assumptions, the richness condition can be seen as a shortcut to relax all such assumptions via the introduction of artificial dominance regions for every action of every player. Formally, this is obtained by embedding the model $(\Theta' \times T')$ in the universal model $(\Theta \times T^*_\Theta)$, where T^*_Θ is the Θ-based universal type space, and Θ is a space that satisfies $\Theta' \subseteq \Theta$ and the richness condition.[12] Then hierarchies of beliefs in $(\Theta' \times T')$ are perturbed by considering sequences in $(\Theta \times T^*_\Theta)$ that converge to the hierarchies of the types in the original model, $(\Theta' \times T')$. This is the spirit of Weinstein and Yildiz's structure theorem.

As discussed in the Introduction, relaxing *all* common knowledge assumptions often entails an unnecessarily demanding robustness test. The structure theorem without richness (Theorem 1) can thus be thought of as a robustness exercise when some (as opposed to all) common knowledge assumptions are relaxed. It allows us to accommodate both cases in which we may wish to maintain common knowledge of a natural parameter space Θ' that does not satisfy richness (e.g., Section 4.1), and cases in which we may want to relax *some*, but not *all* common knowledge assumptions, by introducing artificial dominance regions for only some of the actions (e.g., Section 4.2). (In the end, the difference is only one of interpretation.)

In this section, I discuss some applications that illustrate the versatility of Theorem 1 in addressing applied and theoretical questions. As I show, Theorem 1 implies that in several contexts, minimal perturbations of common knowledge assumptions may produce results as strong as Weinstein and Yildiz's, both in terms of (non-) robustness (Sections 4.1 and 4.2) and in terms of generic uniqueness (Section 4.3). Section 4.4 presents a characterization result, also obtained from minor modifications of the proof of Theorem 1: Proposition 3 shows that the set of actions that are uniquely rationalizable for some hierarchy of beliefs coincides with the set of actions for which a nearby uniqueness result analogous to Theorem 1 holds. Proposition 3 thus characterizes the structure of ICR on arbitrary spaces of uncertainty.

4.1 *Robustness under common knowledge of the rules of the game*

In many settings, players' payoff functions are derived from agents' preferences over an underlying space of outcomes X, and some rules of the game that define agents' actions and an outcome function relating actions to outcomes. Such rules of the game may be given by a properly designed mechanism or by the details of the institutions in which agents' interact (an auction, a bargaining protocol, a market, etc.). Preferences over outcomes are represented by (possibly state-dependent) utility functions $U_i : X \times \Theta \to \mathbb{R}$, while the rules of the game are represented by a tuple $\mathcal{G} = \langle N, (A_i)_{i\in N}, g\rangle$, where $g : A \to X$ is the outcome function. The resulting game G is obtained by setting $u_i = U_i \circ g$. The richness condition is expressed in terms of the payoff functions in G, which combine common knowledge assumptions about agents' preferences with ones about the rules of the game.

[12]The universal model terminology is from Penta (2012a). I refer to that paper for a more thorough discussion of the approach.

In some settings (e.g., in a model of complex market interactions), assuming common knowledge of $\mathcal{G}$ may be as hard to justify as assuming common knowledge of agents' preferences. It may thus be interesting to perturb common knowledge both of preferences and of the rules of the game, and work directly with the derived model G.

In other settings, perturbing common knowledge of the rules of the game may be unreasonable, while it may still be sensible to relax common knowledge of agents' preferences. When this is the case, the richness condition is typically not satisfied and it may be useful to apply Theorem 1 instead. Consider the example presented in the Introduction.

Example 1 (The public good demand game). A society must decide on the quantity of a public good, $x \in X = [0, 1]$. The realization of the public good is delegated to a public authority, which operates according to the following protocol: every agent proposes a quantity and the average of the proposals is implemented.

Agents have single-peaked preferences that depend on the realization of a state of nature, $\theta \in \Theta \subseteq \mathbb{R}$. For simplicity, assume that the society is made of two individuals and that preferences are represented by utility functions $U_i : X \times \Theta \to \mathbb{R}$ such that $U_i(x, \theta) = -(\theta - x)^2$. (This is a pure common value problem, where states $\theta \in \Theta \subseteq \mathbb{R}$ correspond to the (commonly) optimal quantities of x.) The rules of the game in this case are $\mathcal{G} = \langle A_1, A_2, g\rangle$, where $A_1, A_2 \subseteq [0, 1]$ are the sets of possible proposals and the outcome function g is given by $g(a_1, a_2) = (a_1 + a_2)/2$. The payoff functions of the resulting game are thus $u_i : A \times \Theta \to \mathbb{R}$:

$$u_i(a_1, a_2, \theta) = -\left(\theta - \tfrac{1}{2}a_1 + a_2\right)^2.$$

For instance, let us assume that $A_i = \Theta = \{0, \frac{1}{4}, \frac{1}{2}, \frac{3}{4}, 1\}$. It is easy to see that actions $a_i = 0, 1$ are the only ones that admit dominance states in Θ (these are, respectively, $\theta = 0, 1$). We thus have $\mathcal{A}_i^0 = \{0, 1\}$ and the richness condition is not satisfied. Nonetheless, it can be verified that $\mathcal{A}_i^1 = \mathcal{A}_i^\infty = A_i$. Despite the failure of the richness condition, Theorem 1 delivers the following conclusion.

Conclusion 1. For each $t_i \in T^*_{i,\Theta}$ and for each $a_i \in \mathrm{ICR}_i(t_i)$, there exists a sequence $\{t_i^\nu\}_{\nu\in\mathbb{N}} \subseteq T^*_{i,\Theta}$ such that $\lim_{\nu\to\infty} t_i^\nu = t_i$ and, for each $\nu \in \mathbb{N}$, $\mathrm{ICR}_i(t_i^\nu) = \{a_i\}$.

For instance, consider the game $G^{\mathcal{T}}$, where the type space $\mathcal{T} = \langle (T_i, \tau_i)_{i=1,2}\rangle$ is defined as $T_i = \{t_i^\theta : \theta \in \Theta\}$ and

$$\tau_i(t_i^\theta)(t_j^{\theta'}, \theta'') = \begin{cases} 1 & \text{if } \theta' = \theta'' = \theta \\ 0 & \text{otherwise.} \end{cases}$$

(Hence, for each $\theta \in \Theta$, type t_i^θ represents common certainty of θ.) Clearly, $G^{\mathcal{T}}$ has one focal equilibrium in which every type t_i^θ demands quantity $x = \theta$. Nonetheless, even without richness, this prediction is *not* robust to perturbations of higher order beliefs. To see this, it is easy to verify that $\mathrm{ICR}_i(t_i^0) = \{0\}$, $\mathrm{ICR}_i(t_i^{1/4}) = \{0, \frac{1}{4}, \frac{1}{2}\}$, $\mathrm{ICR}_i(t_i^{1/2}) = A_i$, $\mathrm{ICR}_i(t_i^{3/4}) = \{\frac{1}{2}, \frac{3}{4}, 1\}$, and $\mathrm{ICR}_i(t_i^1) = \{1\}$. Then Conclusion 1 implies that, for $\theta \neq 0, 1$, there are types arbitrarily close to t_i^θ for whom the unique ICR (hence equilibrium) action is different from that played by type t_i^θ in the focal equilibrium. $\diamondsuit$

Example 1 involves an extremely mild robustness test, in that only agents' higher order beliefs about θ are perturbed. In particular, by maintaining common knowledge of $\langle \mathcal{G}, \Theta, (U_i)_{i=1,2} \rangle$, we maintained common knowledge of both the rules of the game and the agents' preferences over outcomes.

Clearly, more demanding tests would deliver equally negative results. But as long as common knowledge of the rules of the game is maintained, the richness condition still is *not* satisfied. For instance, let $\Theta^* = (\mathbb{R} \times \mathbb{R})^X$ and $U_i^* : X \times \Theta^* \to \mathbb{R}$ be such that $U_i^*(x, \theta) = \theta_i(x)$. Assuming common knowledge of $\langle \Theta^*, (U_1^*, U_2^*) \rangle$ imposes no restrictions on agents' preferences over X. Yet as long as common knowledge of $\mathcal{G}$ is maintained, the richness condition is *not* satisfied by the payoff functions $u_i = U_i^* \circ g$.

Thus, in the context of Example 1, the richness condition inherently involves either relaxing common knowledge of the rules of the game or that players ultimately care about the final outcome, $x \in X$. A natural objection is that for most applied settings, this kind of robustness test may be unnecessarily demanding. This objection diminishes the practical relevance of Weinstein and Yildiz's results, but is overcome by Theorem 1: in Example 1, Results 1 and 2 hold, despite the failure of the richness condition, even when a very mild robustness test is considered.

4.2 *A structure theorem for equilibrium refinements*

As already discussed, Weinstein and Yildiz's richness condition can be seen as a way to relax common knowledge assumptions via the introduction of artificial dominance regions for every action of every player.

Theorem 1 may be interpreted the same way. Here, the richness of $\mathcal{A}_i^0$ (hence of Θ) captures the extent to which common knowledge assumptions are relaxed, hence the strength of the robustness test: the richer are the $\mathcal{A}_i^0$'s, the more demanding is the robustness test. For instance, if $\mathcal{A}_i^0 = A_i$ for every i (the richness condition), then we are relaxing essentially all common knowledge assumptions. If, instead, we intend to relax only common knowledge that some actions $\hat{A}_i \subseteq A_i$ are not dominant, then we can set $\mathcal{A}_i^0 = \hat{A}_i$.

In many situations, setting $\mathcal{A}_i^0 = A_i$ entails an unreasonably demanding robustness test. Consider auction settings for instance. Depending on the rules of the auction and on the assumptions on bidders' valuations, there may be bids that are dominated for all types of players. Assuming $\mathcal{A}_i^0 = A_i$ for all i introduces dominance regions for bids that are inconsistent with common knowledge of the rules of the auction, or that players are rational, or that payoffs are decreasing in the price paid, and so forth. A much weaker robustness test would be, for instance, to assume that $\mathcal{A}_i^0$ only contains player i's rationalizable bids, i.e., those bids that are consistent with common knowledge of the rules of the game, of players' rationality, etc. Despite the significantly weaker robustness test, point 3 in Section 3.2 implies that Weinstein and Yildiz's nonrobustness results still hold in their full strength.

Example 2. Consider two agents competing in a second-price auction for the allocation of one unit of an indivisible good. Assume that agents' possible bids and types are from a finite grid. For $m \in \mathbb{N}$, let $d \equiv \frac{1}{m}$. Agents' types are denoted by t_i and belong to the set $T_i = \{d, 2d, \ldots, 1-d, 1\}$; the set of possible bids is $A_i = \{0, \frac{d}{n}, \frac{2d}{n}, \ldots, M - \frac{d}{n}, M\}$, where n and M are "large" natural numbers.[13] Let $\theta \in \Theta$ denote the value of the good and assume that agents are risk neutral. Type t_i's beliefs on $\Theta \times T_{-i}$ are such that for every $\hat{t}_{-i} \in T_{-i}$ and $\hat{\theta} = t_i + \hat{t}_{-i}$, $\tau_i(t_i)(\hat{\theta}, \hat{t}_{-i}) = \frac{1}{m}$. In other words, conditional on knowing the opponent's type $\hat{t}_{-i}$, type t_i believes that the value of the object is $t_i + \hat{t}_{-i}$.[14] In the unique symmetric equilibrium, each player bids twice his type, but it is easy to see that this game has many asymmetric equilibria. For instance, for any $k \in \mathbb{N}$, bidding strategies $s_i^*(t_i) = (1+kn)t_i$ and $s_j^*(t_j) = (1 + \frac{1}{kn})t_j$ form a Bayesian equilibrium.[15] In fact, it can be shown that all bids greater than or equal to $t_i + d$ are rationalizable for type t_i.

One implication of Weinstein and Yildiz's results in this example is that each such tuple of equilibrium strategies is uniquely rationalizable for some arbitrarily close model of beliefs. In the present framework, this is obtained by setting $\mathcal{A}_i^0 = A_i$, which allows bids that are weakly dominated for all types (e.g., bidding less than d) to be strictly dominant in some state of the world.

A considerably weaker robustness test is to allow dominance regions *only for the equilibrium bids*; that is, let $E_i(t_i) \subseteq A_i$ denote the set of *equilibrium bids* for type t_i and set $\mathcal{A}_i^0 = \bigcup_{t_i \in T_i} E_i(t_i)$. Then any such equilibrium bid is rationalizable (clearly, $E_i(t_i) \subseteq \mathrm{ICR}_i(t_i)$ for all t_i) and (by definition of equilibrium) it is justified by conjectures concentrated on the opponents' equilibrium bids, $E_j(t_j) \subseteq \mathrm{ICR}_j(t_j)$. But since $E_j(t_j) \subseteq \mathcal{A}_j^0 \subseteq \mathcal{A}_j^\infty$, the conditions of Theorem 1 are satisfied (that is, $E_i(t_i) \subseteq \mathrm{ICR}_i(t_i; \mathcal{A})$). Hence we conclude that for every equilibrium bid $a_i^* \in E_i(t_i)$, there exists a sequence $\{t_i^\nu\}_{\nu \in \mathbb{N}}$ converging to t_i such that for each $\nu \in \mathbb{N}$, $\{a_i^*\} = \mathrm{ICR}_i(t_i^\nu)$ (hence also $\{a_i^*\} = E_i(t_i^\nu)$). We thus obtain the following result: any equilibrium bid is uniquely rationalizable for an arbitrarily close model of beliefs. ◊

The next general result follows immediately from the logic of Example 2 and from Theorem 1.

Proposition 1. *Fix a Bayesian game $G^{\mathcal{T}}$. For each $i \in I$ and $t_i \in T_i$, let $E_i(t_i) \subseteq A_i$ denote the set of* equilibrium actions *for type t_i. If, for every i, $\mathcal{A}_i^0 = \bigcup_{t_i \in T_i} E_i(t_i)$, then for every i, for every t_i, and for every $a_i \in E_i(t_i)$, there exists a sequence $\{t_i^\nu\} \subseteq \hat{T}_i$ s.t. $t_i^\nu \to t_i$, and for each $\nu \in \mathbb{N}$, $\{a_i\} = \mathrm{ICR}_i(t_i^\nu) = E_i(t_i^\nu)$.*

[13]For illustrating purposes, it is convenient to use a finer grid for the bids than for the types, hence n. Similarly, the upper bound M is needed simply to maintain finiteness of the actions space, but can be set arbitrarily large. Alternatively, M can be thought of as a maximum liability constraint for the agent.

[14]This example is simply a discrete version of a *wallet game* with independently and uniformly distributed types (e.g., Klemperer 1998). These games are usually modelled with no reference to θ, with the values equal to $t_1 + t_2$. The presence of θ here plays no role other than to facilitate the comparison with the setup of the previous section.

[15]The finer grid for the bids is assumed to guarantee that these strategies take values in A_i. In the absence of such a finer grid, it could be that (e.g.) for some t_j, $(1 + 1/(kn))t_j$ falls in between two points in the grid. In that case, one of the adjacent bids is the best reply. The calculations are more involved, without changing the main insight.

Remark 3. Notice that the result in Proposition 1 remains true if for every player and type t_i, $E_i(t_i)$ represents the set of actions of type t_i that are consistent with an arbitrary equilibrium refinement E; for instance, the set of actions played in non-weakly-dominated equilibria, or in trembling hand perfect equilibria, or the set of actions played in monotone strategy equilibria, and so on. The key for the result of Proposition 1 is that actions in $E_i(t_i)$ must be justified by conjectures concentrated on $E_{-i}(t_{-i})$, so that the condition $E_i(t_i) \subseteq \text{ICR}_i(t_i; \mathcal{A})$ is satisfied if $\mathcal{A}_j^0 = \bigcup_{t_j \in T_j} E_j(t_j)$ for all $j \in I$.

Behavioral types and equilibrium strategies In this subsection, I discuss an alternative interpretation of Weinstein and Yildiz's result (and of the results above).

By introducing artificial dominance regions for every action, the richness condition implies the existence, in the universal type space, of types that play a specific action independently of their beliefs about the opponents' behavior. The analogous objects in Theorem 1 are the types $t_i \in \bar{T}_i$ introduced in Definition 2.

These dominance types play a role similar to that of the commitment types in the reputation literature, or of the noise traders sometimes used in auctions or in the literature on strategic foundations of the rational expectations equilibrium (REE).[16] From this viewpoint, Weinstein and Yildiz's result (Result 1) can be rephrased as follows: if behavioral types for all actions are introduced, then no refinement of ICR delivers predictions that are robust to perturbations of higher order beliefs. Theorem 1 can be rephrased in a similar way, once the sets $\mathcal{A}_i^0$ are interpreted as the sets of actions for which behavioral types are introduced. Once again, the richness of $\mathcal{A}_i^0$ captures the strength of the robustness exercise: the richer is the set of behavioral types, the more demanding is the robustness test.

In applied models that adopt behavioral types, it is often the case that such types play irrationally (these types are also called crazy types). There is a sense in which the introduction of irrational behavioral types constitutes a stronger perturbation of the benchmark than the introduction of types committed to rational or equilibrium behavior. When interpreted in terms of behavioral types, Proposition 1 delivers a nonrobustness result for arbitrary equilibrium concepts E, under the minimal perturbation that higher order beliefs simply do not rule out that the opponents may be committed to some specific behavior consistent with E itself.[17]

[16] For the literature on reputation, see Chapters 15–18 in Mailath and Samuelson (2006). In the REE literature, the seminal paper on noise traders is Kyle (1985). (For more on noise traders, see, e.g., Dow and Gorton 2008 and the references therein.) In the context of auctions, Cho and Lee (2011) consider noise bidders in a first-price auction, adopting ICR as the solution concept.

[17] When the benchmark model is a nondegenerate Bayesian game (i.e., such that the sets of types are nonsingletons), equilibrium strategies are functions $\sigma_i : T_i \to A_i$, not actions $a_i \in A_i$. When behavioral types are introduced in these models, they are typically modelled as types that play some specific strategy $\sigma_i^* : T_i \to A_i$. This can be mapped to the model above as enriching the sets T_i by adding, for each $t_i \in T_i$, a "behavioral replica" t_i^c committed to playing $a_i = \sigma_i^*(t_i)$. A model in which i may be committed to the strategy σ_i^* can thus be thought of as setting $\mathcal{A}_i^0 = \bigcup_{t_i \in T_i} \sigma_i^*(t_i)$; that is, a robustness exercise in which higher order beliefs do not rule out the replica types $\{t_i^c : t_i \in T_i\}$ who play $\sigma_i^*(t_i)$.

4.3 *Generic uniqueness without richness*

Multiplicity is pervasive in applied models. Weinstein and Yildiz's (2007) generic uniqueness result (Result 2) can be interpreted as saying that the typical indeterminacy of standard game theoretic models is often the consequence of the common knowledge assumptions implicit in game theoretic models. If such assumptions are relaxed, hierarchies of beliefs "typically" have a unique rationalizable outcome.

Result 2 thus generalizes the main insight from the literature on global games. At a formal level, though, the two approaches are quite different and difficult to compare. Subsequently, I consider an application of Theorem 1 that shows how very weak relaxations of common knowledge assumptions may suffice to guarantee the full version of Weinstein and Yildiz's results. This result is used to shed some light on the connection between Weinstein and Yildiz's and the global games approach.

PROPOSITION 2. *Suppose that Θ is such that the following statements hold.*

(i) For each i, there exist dominance states $\theta^{a_i'}, \theta^{a_i''} \in \Theta$ for some distinct $a_i', a_i'' \in A_i$.

(ii) There exists $\hat{\theta} \in \Theta$ such that for every i and for every $\hat{a}_i \in A_i$, $\exists p^{\hat{a}_i} \in [0, 1]$,

$$\{\hat{a}_i\} = \arg\max_{a_i \in A_i} p^{\hat{a}_i} u_i(a_i, a_{-i}'; \hat{\theta}) + (1 - p^{\hat{a}_i}) u_i(a_i, a_{-i}''; \hat{\theta}). \tag{4}$$

Then $\mathcal{A}_i^{\infty} = A_i$ for every i.

PROOF. Condition (i) implies that $\mathcal{A}_i^0 \supseteq \{a_i', a_i''\}$. By definition,

$$\mathcal{A}_i^1 := \{a_i \in A_i : \exists \beta^i \in \Delta(\Theta \times \mathcal{A}_{-i}^0) \text{ s.t. } \{a_i\} = \mathrm{BR}_i(\beta^i)\}.$$

That $\mathcal{A}_i^1 = A_i$ (hence $\mathcal{A}_i^{\infty} = A_i$) follows from condition (ii) simply by letting $\beta^i \in \Delta(\Theta \times \mathcal{A}_{-i}^0)$ be such that $\beta^i(\{\hat{\theta}\} \times \{a_{-j}'\}) = p^{a_i}$ from every a_i. □

COROLLARY 1. *If Θ satisfies conditions (i) and (ii) in Proposition 2, Weinstein and Yildiz's (2007) results (Results 1 and 2) obtain.*

Notice that Example 1 provides one instance in which the conditions of Proposition 2 are satisfied: actions $a_i = 0, 1$ admits dominance regions, respectively, at states $\theta = 0, 1$ (condition (i)), while condition (ii) is satisfied at $\hat{\theta} = 1/2$.

In general, Proposition 2 requires only the existence of dominance regions for two actions of every player, a_i' and a_i'', plus that of a payoff state $\hat{\theta}$ in which every action of every player is a strict best response to conjectures concentrated on a_{-i}' and a_{-i}''. Hence, if Θ satisfies the conditions of Proposition 2, then it is *not* common knowledge that a_i' and a_i'' are *not* dominant, and it is not common knowledge that payoffs *do not* satisfy condition (ii).

Several restrictions on $u_i(\cdot, \hat{\theta})$ may be provided so that condition (ii) is satisfied. I next consider one in particular that shows how Proposition 2 can be used to connect the approaches used by Weinstein and Yildiz and in the global games literature.

Similar to the global games literature, condition (i) in Proposition 2 requires only the existence of dominance regions for two actions of each player. (In the global games literature, these two actions are typically taken to be the highest and lowest actions according to some ordering of the action space.) In addition to this, the global games literature assumes that players' actions are strategic complements at all states (plus other technical assumptions); that is, common knowledge of strategic complementarity is assumed.

Lemma 5 below shows that condition (ii) is satisfied if payoff functions at state $\hat{\theta}$ exhibit strategic complementarities, plus some conditions interpretable as strict concavity in own action. This way, Proposition 2 can be used to obtain a generic uniqueness result based on the existence of dominance regions for only two actions of each player, but without assuming common knowledge of strategic complementarity: all that is required is that it is *not* common knowledge that actions are *not* strategic complements.

Noncommon knowledge of noncomplementarity Let $(\lhd^i, X_i)_{i\in I}$ be a collection of completely asymmetrically ordered sets ($\lhd^i$ is a complete, transitive, irreflexive, and antisymmetric relation on X_i). I denote by $\unlhd^{-i}$ the natural partial order on $X_{-i} = \times_{j\neq i} X_j$ induced by the orders $\lhd^j$, and I refer to its symmetric and asymmetric parts, respectively, as $=^{-i}$ and $\lhd^{-i}$. (Clearly, $(\unlhd^{-i}, X_{-i})$ and $(\unlhd, X)$ are complete lattices.) If X_i is finite, it is convenient to represent order $\lhd^i$ by properly labelling elements $x_i \in X_i$, so that $x_i^0 \lhd^i x_i^1 \lhd^i \cdots \lhd^i x_i^{|X_i|-1}$. Real function $g: X \to \mathbb{R}$ is *strictly* $\lhd^{-i}$*-increasing* in $x_{-i} \in X_{-i}$ if for every x_i, $g(x_i, x_{-i}) > g(x_i, x'_{-i})$ whenever $x'_{-i} \lhd^{-i} x_{-i}$.

Next, I introduce assumptions on $\hat{\theta}$ that guarantee that (4) in condition (ii) of Proposition 2 is satisfied.

Assumptions. Let $\hat{\theta} \in \Theta$ be such that for every i, the following statements hold.

A.1 *No dominated actions.* For all $a_i \in A_i$, $\exists \psi^i \in \Delta(\{\hat{\theta}\} \times A_{-i}) : a_i \in \mathrm{BR}_i(\psi^i)$.

A.2 *Strategic complementaries.* There exist complete orders on each A_i, $a^0 \lhd^i a_i^1 \lhd^i \cdots \lhd^i a_i^{N_i}$ (where $N_i = |A_i| - 1$) such that for every i, for each $m \in \{1, \dots, N_i\}$ and $n \in \{0, \dots, m-1\}$, $u_i(a_i^m, a_{-i}; \hat{\theta}) - u_i(a_i^n, a_{-i}; \hat{\theta})$ is $\unlhd^{-i}$-strictly increasing in a_{-i}.

A.3 *"Diminishing increments."*

(a) For each $a_{-i} \in A_{-i}$ and $n \in \{1, \dots, N_i\}$,

$$\Delta u_i(n, a_{-i}) \equiv u_i(a_i^n, a_{-i}; \hat{\theta}) - u_i(a_i^{n-1}, a_{-i}; \hat{\theta})$$

is *strictly decreasing in* n.

(b) For each $m \in \{1, \dots, N_i\}$ and $n \in \{0, \dots, m-1\}$,

$$\frac{\Delta u_i(n, a_{-i}^N)}{u_i(a_i^m, a_{-i}^{N_{-i}}) - u_i(a_i^n, a_{-i}^{N_{-i}})} > \frac{\Delta u_i(n, a_{-i}^0)}{u_i(a_i^m, a_{-i}^0) - u_i(a_i^n, a_{-i}^0)}$$

(where $a_{-i}^0 = (a_j^0)_{j\neq i}$ and $a_{-i}^{N_{-i}} = (a_j^{N_j})_{j\neq i}$).

Assumption A.1 requires that no action is strictly dominated at $\hat{\theta}$. Assumption A.2 says that for some order on player' actions, each player's incentive to raise her action at $\hat{\theta}$ is strictly increasing in her opponents' actions. Assumption A.3(a) is a discrete version of a strict concavity in own action assumption. It requires that payoff increments from increasing own action are strictly decreasing for any action profile of the opponents. Assumption A.3(b) restricts the relative payoff increment from own action at the opponents' highest and lowest action profiles. It requires that the ratio between the marginal increment $\Delta u_i(n, a_{-i}^N)$ and the increment moving from action a_i^n to a higher action a_i^m is larger when the opponents play the highest action profile than when they play the lowest action profile. This assumption guarantees that each player's preferences over own actions are single-peaked with respect to the probability weight assigned to the opponents' highest action profile, $a_{-i}^{N_{-i}}$ (when complementary probability is assigned to the lowest action profile, a_{-i}^0).

Lemma 5. *If $\hat{\theta}$ satisfies Assumptions A.1–A.3, then condition (ii) in Proposition 2 is satisfied.*

For the proof, see the Appendix.

4.4 *Nonrobustness on arbitrary spaces of uncertainty*

In this section, I present a result similar to Theorem 1 that reveals something more on the structure of rationalizability on arbitrary spaces of uncertainty. Fix Θ and consider the solution concept correspondence $\text{ICR}_i : T_i^* \rightrightarrows A_i$ that assigns to each hierarchy of beliefs the corresponding set of rationalizable actions for player i. Define $\mathcal{U}_i \subseteq A_i$ as the set of player i's actions that are uniquely rationalizable for some type in the universal type space, and let $\mathcal{A}^* = \times_{i\in N}\mathcal{A}_i^*$ s.t. $\mathcal{A}_i^* \subseteq \mathcal{U}_i$ for each i denote the largest subset of $\mathcal{U} = \times_{i\in N}\mathcal{U}_i$ that satisfies the property

$$\mathcal{A}_i^* = \{a_i \in \mathcal{U}_i : \exists \varphi \in \Delta(\Theta \times \mathcal{A}_{-i}^*) \text{ s.t. } \text{BR}_i(\varphi) = \{a_i\}\}.$$

Before moving to the result of this section, consider the following observations.

Remark 4. Consider the following observations.

(i) If $\mathcal{A}_i^0$ is set equal to $\mathcal{A}_i^*$ for every i and the recursion (1) is applied, then we obtain $\mathcal{A}_i^0 = \mathcal{A}_i^* = \mathcal{A}_i^k$ for every k; hence, $\mathcal{A}_i^0 = \mathcal{A}_i^* = \mathcal{A}_i^\infty$. This implies that $\mathcal{W}_i^k \subseteq \mathcal{W}_i^{k+1}$ for all $k \in \mathbb{N}$ (Definition 3), not just for all $k \geq K$.

(ii) For any $a_i \in \mathcal{U}_i$, there exists a finite type $t_i^{a_i} \in \hat{T}_i$ such that $\text{ICR}_i(t_i^{a_i}) = \{a_i\}$. This follows from the definition of $\mathcal{U}_i$, which implies that $\exists t_i \in T_i^*$ s.t. $\text{ICR}_i(t_i) = \{a_i\}$, and the fact that ICR_i is upper hemicontinuous on T_i^* (Dekel et al. 2006) and $\hat{T}_i$ is dense in T_i^* (Mertens and Zamir 1985).

The next result shows that the set $\mathcal{A}^*$ characterizes the set of actions for which a structure theorem analogous to Theorem 1 holds: for each t_i such that $a_i \in \text{ICR}_i(t_i) \cap \mathcal{A}_i^*$

that is justified by conjectures concentrated on $\mathcal{A}^*_{-i}$, we can construct a sequence of (finite) types converging to t_i for which a_i is uniquely rationalizable.

Formally, let

$$\mathrm{ICR}_i(t_i;\mathcal{A}^*) = \{a_i \in \mathrm{ICR}_i(t_i) \cap \mathcal{A}^*_i : \exists \psi^{a_i} \in \Psi_i(t_i) \text{ s.t. } a_i \in \mathrm{BR}_i(\psi^{a_i}) \text{ and } \operatorname{supp}(\operatorname{marg}_{A_{-i}} \psi^{a_i}) \subseteq \mathcal{A}^*_{-i}\}.$$

Since, by definition, $\mathrm{ICR}_i(t_i;\mathcal{A}^*) \subseteq \mathcal{A}^*_i \subseteq \mathcal{U}_i$, we already know from Remark 4(ii) that any $a_i \in \mathrm{ICR}_i(t_i;\mathcal{A}^*)$ is uniquely rationalizable for some finite type $t_i^{a_i} \in \hat{T}_i$. Such $t_i^{a_i}$, however, may be very far from t_i. The next proposition shows that, in fact, a_i is also uniquely rationalizable for some type arbitrarily close to t_i.

Proposition 3. *Fix an arbitrary space of uncertainty, Θ. For every $\hat{t}_i \in \hat{T}_i$ and $\hat{a}_i \in \mathrm{ICR}_i(\hat{t}_i,\mathcal{A}^*)$, there exists a sequence $\{t_i^\nu\} \subseteq \hat{T}_i$ such that $t_i^\nu \to \hat{t}_i$ and $\{\hat{a}_i\} = \mathrm{ICR}_i(t_i^\nu)$ for each $\nu \in \mathbb{N}$.*

Proof. The result follows from minor adaptations of the proof of Theorem 1. The only changes are required to prove the following step.

Step 1. *If $a_i \in \mathrm{ICR}_i(t_i,\mathcal{A}^*)$, then there exists $t_i(\varepsilon) \to t_i$ as $\varepsilon \to 0$ such that for each $\varepsilon > 0$, $a_i \in \mathcal{W}_i(t_i(\varepsilon))$ and $t_i(\varepsilon) \in \hat{T}_i$ (hence, $a_i \in \mathcal{W}_i^k(t_i(\varepsilon))$ for all k).*

This step is analogous to Lemma 3 above. The statement does not require $k \geq K$ because of Remark 4(i). Given Step 1, the rest of the proof follows as in Theorem 1.

Proof of Step 1. Using Remark 4(ii), and the fact that $\mathcal{A}^*_i \subseteq \mathcal{U}_i$, define the set $\tilde{T}_i = \{t_i^{a_i} : a_i \in \mathcal{A}^*_i\} \subseteq \hat{T}_i$. Notice that because of the definition of $\mathcal{A}^*$, such types $t_i^a \in \tilde{T}_i$ can be chosen so that $\tilde{T} = \times_{i\in N}\tilde{T}_i$ is a belief-closed subset of $\hat{T}$. Hence, there exists a type space $(\tilde{T}_i, \tilde{\tau}_i)_{i\in N}$ such that $\mathrm{ICR}_i(t_i^{a_i}) = \{a_i\}$ for each $t_i^{a_i} \in \tilde{T}_i$. For each i, define $\mathrm{ICR}_i^* = \{(t_i, a_i) : a_i \in \mathrm{ICR}_i(t_i;\mathcal{A}_i^*)\}$. The rest of the proof is the same as in Lemma 3, once the sets $\mathrm{ICR}_i^{\mathcal{A}}$, $\mathcal{A}_i^\infty$, and $\bar{T}_i$ are replaced, respectively, by the sets ICR_i^*, $\mathcal{A}_i^*$, and $\tilde{T}_i$, and the beliefs in (3) are replaced by

$$\tau_i^\varepsilon(\bar{\iota}_i^\varepsilon(t_i, a_i)) = \varepsilon \cdot \tilde{\tau}_i(t_i^{a_i}) + (1-\varepsilon)[\psi^{a_i} \circ \hat{\iota}^{-1}_{-i,\varepsilon}].$$

(Notice that for each $t_i^{a_i} \in \tilde{T}_i$, beliefs $\tilde{\tau}_i(t_i^{a_i})$ are such that a_i is uniquely rationalizable. Hence $\{a_i\} = \mathrm{BR}_i(\beta_i)$ for any conjecture β_i that is rationalizable for type $t_i^{a_i}$.) □

5. Discussion

Propositions 1, 2, and 3 are only some of the implications of Theorem 1. As these propositions show, Theorem 1 can be easily applied to a variety of theoretical and applied problems. In the following subsections, I discuss the related literature and some possible directions for future research based on the results presented herein.

Related literature

Weinstein and Yildiz (2007) prove the first structure theorem for ICR in static finite games, assuming the richness condition on Θ. Their approach cannot accommodate dynamic games because the richness condition is not satisfied in dynamic games. Structure theorems for dynamic games are provided by Chen (2012) and Penta (2012a).

Chen (2012) modifies the notion of richness so as to extend the structure theorem for ICR to finite dynamic games in normal form. Penta (2012a) instead observes that an implicit assumption of Weinstein and Yildiz is that players have no information about payoffs, and addresses the question of robustness to perturbations of higher order beliefs under arbitrary information structures. Penta (2012a) characterizes the "robust predictions" in static and dynamic (finite) games under arbitrary information structures. This characterization is provided by an extensive form solution concept—interim sequential rationalizability (ISR). In static games, ISR coincides with ICR and it does not depend on the assumptions on agents' information. This is not the case in dynamic games, where ISR refines ICR and depends on the details of the information structure. The no information assumption therefore entails no loss of generality in static settings, but does not hold in dynamic settings.

Recently, Weinstein and Yildiz (2012) extend these results to infinite horizon games with payoffs continuous at infinity. Weinstein and Yildiz (2011) instead address the robustness of equilibrium behavior in nice (static) games.

Unlike the present paper, Weinstein and Yildiz (2007), Chen (2012), and Penta (2012a) analyze the extreme case in which all common knowledge assumptions are relaxed, in the sense that they prove structure theorems under various formalizations of the richness condition.[18] It should not be difficult to extend Theorem 1 to games with a continuum of actions, adopting conditions similar to those considered by Weinstein and Yildiz (2011).

Common knowledge of payoffs and mechanism design

In a standard mechanism design problem, the primitives of the environment are given by a set of outcomes X, agents' preferences $u_i : X \times \Theta \to \mathbb{R}$, and agents' information $t_i \in T_i$ about the payoff state and the opponents' information, (θ, t_{-i}). Typically, payoff functions and agents' beliefs are assumed to be common knowledge. In these contexts, one might be concerned with the robustness of implementation results when such common knowledge assumptions are relaxed.[19] In such mechanism design problems, the richness of Θ cannot be expressed in terms of dominance regions for players' actions,

[18]Weinstein and Yildiz (2011) also pursue a relaxation of the richness assumption, but they do so by adopting a stronger solution concept than rationalizability—Bayesian equilibrium. The stronger solution concept allows them to obtain stronger results without requiring the existence of dominance regions. Similarly, for infinitely repeated games, Weinstein and Yildiz (2012) also provide versions of the structure theorem that maintain common knowledge of some of the structure of the game (namely, the fact that the game being played is a repeated game).

[19]The exercise that I consider here differs from that considered by recent literature on *robust mechanism design* (e.g., Bergemann and Morris 2005, 2009 or Penta 2012b for dynamic environments). Those papers focus on situations in which nothing is known about agents' beliefs, but payoff functions are maintained as

because such actions and the mapping from action profiles to the space of outcomes is chosen by the mechanism designer. Hence, the set of actions for which dominance regions exist may vary with the mechanism chosen. In these cases, Theorem 1 may be used instead: for given assumptions on Θ and for given mechanism $\mathcal{M}$, one can compute the set $\mathcal{A}_i^{0,\mathcal{M}}$ of actions that are uniquely rationalizable for some type *given mechanism* $\mathcal{M}$ and then apply Theorem 1 to the rationalizable actions in any given mechanism.[20]

Extensive-form robustness

More generally, unlike results based on the richness condition, Theorem 1 allows studying situations in which some assumptions are maintained as common knowledge, while others are relaxed. In the mechanism design problems discussed above, the implicit assumption is that while common knowledge assumptions on payoffs are relaxed, the rules of the mechanism are maintained as common knowledge. One might imagine, instead, situations in which, given certain common knowledge restrictions on payoffs, there may be higher order uncertainty concerning the *rules of the game*. This is an important theoretical question, which is thoroughly neglected by the literature. Theorem 1 can be adapted to study these situations as well: given a tuple $\mathcal{E} = \langle I, X, \Theta, (T_i, \tau_i, u_i)_{i\in I}\rangle$ (assumed common knowledge) that represents agents' preferences over outcome space X and their beliefs, and letting A denote the set of action profiles available to the agents, uncertainty about the rules of the game can be represented by considering a set $\mathcal{O}$ of possible outcome functions, $O\colon A \to X$. Given tuple $\mathcal{E}$, for any set $\mathcal{O}$ representing the assumptions on the rules of the game that are maintained as common knowledge, one may compute the set $\mathcal{A}_i^{0,\mathcal{O}}$ of actions that are uniquely rationalizable for some beliefs over $\mathcal{O}$, and accordingly apply Theorem 1.

Appendix

Proof of Lemma 5. For any $p \in [0,1]$, let $\mathrm{BR}_i(p) = \arg\max_{a_i\in A_i} p u_i(a_i, a_{-i}^{N_{-i}}; \hat{\theta}) + (1-p)u_i(a_i, a_{-i}^0; \hat{\theta})$. From Assumptions A.1 and A.2, it follows that $a_i^0 \in \mathrm{BR}_i(0)$ and $a_i^N \in \mathrm{BR}_i(1)$, i.e., $\Delta u_i(1, a_{-i}^0) \leq 0$ and $\Delta u_i(N_i, a_{-i}^{N_{-i}}) \geq 0$. Then A.3(a) implies that $\Delta u_i(n, a_{-i}^0) < 0$ for every $n > 1$ and $\Delta u_i(n, a_{-i}^{N_{-i}}) > 0$ for every $n < N_i$.

Now let p^{i,N_i} be the probability weight on $a_{-i}^{N_{-i}}$ (the complementary weight being on a_{-i}^0) that makes i indifferent between action $a_i^{N_i}$ and a_i^n; that is, $p^{i,N}(n)$ solves

$$p^{i,N_i}(n)[u_i(a_i^{N_i}, a_{-i}^{N_{-i}}; \hat{\theta})] + (1-p^{i,N_i}(n))[u_i(a_i^{N_i}, a_{-i}^0; \hat{\theta})] = p^{i,N_i}(n)[u_i(a_i^n, a_{-i}^{N_{-i}}; \hat{\theta})] + (1-p^{i,N_i}(n))[u_i(a_i^n, a_{-i}^0; \hat{\theta})],$$

common knowledge. Here, instead, I am considering a robustness exercise in which beliefs are maintained close to some benchmark, but higher order beliefs may not rule out significantly different payoffs. Hence, payoffs are not common knowledge.

[20] Oury and Tercieux (2012) recently applies Weinstein and Yildiz's results to problems of mechanism design, introducing states in which the messages in a given mechanism are payoff-relevant. Dispensing with the richness condition, Theorem 1 may be used to improve on those results by maintaining the assumption that messages are payoff-irrelevant.

hence,

$$p^{i,N_i}(n) = \frac{-\sum_{k=n+1}^{N_i} \Delta u_i(k, a_{-i}^0)}{(\sum_{k=n+1}^{N_i} \Delta u_i(k, a_{-i}^{N_{-i}}) - \Delta u_i(k, a_{-i}^0))}.$$

(Notice that this is a well defined probability: since $\Delta u_i(n, a_{-i}^0) < 0$ for every $n > 1$ and $\Delta u_i(1, a_{-i}^0) \leq 0$, the numerator is positive. Assumption A.2 guarantees that $\Delta u_i(k, a_{-i})$ is increasing in a_{-i} for every k, hence $\sum_{k=n+1}^{N_i} \Delta u_i(k, a_{-i}^{N_{-i}}) > \sum_{k=n+1}^{N_i} \Delta u_i(k, a_{-i}^0)$.)

Finally, A.3(b) guarantees that $p^{i,N_i}(n)$ is decreasing in n:

$$p^{i,N_i}(n) > p^{i,N_i}(n-1)$$

if and only if

$$\frac{\sum_{k=n+1}^{N_i} \Delta u_i(k, a_{-i}^0)}{(\sum_{k=n+1}^{N_i} \Delta u_i(k, a_{-i}^{N_{-i}}) - \Delta u_i(k, a_{-i}^0))} < \frac{\sum_{k=n}^{N_i} \Delta u_i(k, a_{-i}^0)}{(\sum_{k=n}^{N_i} \Delta u_i(k, a_{-i}^{N_{-i}}) - \Delta u_i(k, a_{-i}^0))}$$

$$\frac{(\sum_{k=n}^{N_i} \Delta u_i(k, a_{-i}^{N_{-i}}) - \Delta u_i(k, a_{-i}^0))}{(\sum_{k=n+1}^{N_i} \Delta u_i(k, a_{-i}^{N_{-i}}) - \Delta u_i(k, a_{-i}^0))} > \frac{\sum_{k=n}^{N_i} \Delta u_i(k, a_{-i}^0)}{\sum_{k=n+1}^{N_i} \Delta u_i(k, a_{-i}^0)}$$

$$\frac{\Delta u_i(n, a_{-i}^{N_{-i}}) - \Delta u_i(n, a_{-i}^0)}{(\sum_{k=n+1}^{N_i} \Delta u_i(k, a_{-i}^{N_{-i}}) - \Delta u_i(k, a_{-i}^0))} > \frac{\Delta u_i(n, a_{-i}^0)}{\sum_{k=n+1}^{N_i} \Delta u_i(k, a_{-i}^0)}$$

$$\frac{\Delta u_i(n, a_{-i}^{N_{-i}}) - \Delta u_i(n, a_{-i}^0)}{\Delta u_i(n, a_{-i}^0)} < \frac{(\sum_{k=n+1}^{N_i} \Delta u_i(k, a_{-i}^{N_{-i}}) - \Delta u_i(k, a_{-i}^0))}{\sum_{k=n+1}^{N_i} \Delta u_i(k, a_{-i}^0)}$$

$$\frac{\Delta u_i(n, a_{-i}^{N_{-i}})}{\Delta u_i(n, a_{-i}^0)} < \frac{\sum_{k=n+1}^{N_i} \Delta u_i(k, a_{-i}^{N_{-i}})}{\sum_{k=n+1}^{N_i} \Delta u_i(k, a_{-i}^0)}$$

$$\frac{\Delta u_i(n, a_{-i}^{N_{-i}})}{\Delta u_i(n, a_{-i}^0)} < \frac{u_i(a_i^{N_i}, a_{-i}^{N_{-i}}) - u_i(n, a_{-i}^{N_{-i}})}{u_i(a_i^{N_i}, a_{-i}^0) - u_i(n, a_{-i}^0)},$$

which is satisfied under condition A.3(b), setting $m = N_i$ in A.3(b). Hence $\{a_i^{N_i-1}\} = \text{BR}_i(p^{i,N_i}(N_i - 1))$ and for any $p < p^{i,N_i}(N_i - 1)$, action $a_i^{N_i-1}$ is strictly preferred to $a_i^{N_i}$.

Now, recursively, for every $m < N_i$, and $n < m$, let $p^{i,m}(n)$ solve

$$\begin{aligned} p^{i,m}(n)[u_i(a_i^m, a_{-i}^{N_{-i}}; \hat{\theta})] + (1 - p^{i,m}(n))[u_i(a_i^m, a_{-i}^0; \hat{\theta})] \\ = p^{i,m}(n)[u_i(a_i^n, a_{-i}^{N_{-i}}; \hat{\theta})] + (1 - p^{i,m}(n))[u_i(a_i^n, a_{-i}^0; \hat{\theta})], \end{aligned}$$

hence,

$$p^{i,m}(n) = \frac{-\sum_{k=n+1}^{m} \Delta u_i(k, a_{-i}^0)}{(\sum_{k=n+1}^{m} \Delta u_i(k, a_{-i}^{N_{-i}}) - \Delta u_i(k, a_{-i}^0))}.$$

Similar to above, it can be shown that A.3(b) guarantees that $p^{i,m}(n)$ is strictly decreasing in n. Additionally, $\{a_i^m\} = \text{BR}_i(p^{i,m}(m))$. Now, for every $m = 1, \dots, N_i - 1$, let $\psi^m = p^{i,m+1}(m)$, $\psi^0 = 0$ and $\psi^{N_i} = 1$. Then for every $n = 0, 1, \ldots, N_i$, we have that $\{a_i^n\} = \mathrm{BR}_i(\psi^n)$. □

References

Battigalli, Pierpaolo, Alfredo Di Tillio, Edoardo Grillo, and Antonio Penta (2011), "Interactive epistemology and solution concepts for games with asymmetric information." *The B.E. Journal of Theoretical Economics*, 11 (1). [411]

Bergemann, Dirk and Stephen Morris (2005), "Robust mechanism design." *Econometrica*, 73, 1771–1813. [426]

Bergemann, Dirk and Stephen Morris (2009), "Robust implementation in direct mechanisms." *Review of Economic Studies*, 76, 1175–1204. [426]

Bernheim, B. Douglas (1984), "Rationalizable strategic behavior." *Econometrica*, 52, 1007–1028. [411]

Chen, Yi-Chun (2012), "A structure theorem for rationalizability in the normal form of dynamic games." *Games and Economic Behavior*, 75, 587–597. [426]

Cho, In-Koo and Soonho Lee (2011), "Rationalizability and monotonicity in large auctions with a single delivery price and noise bidders." Unpublished paper. [421]

Dekel, Eddie, Drew Fudenberg, and Stephen Morris (2006), "Topologies on types." *Theoretical Economics*, 1, 275–309. [424]

Dekel, Eddie, Drew Fudenberg, and Stephen Morris (2007), "Interim correlated rationalizability." *Theoretical Economics*, 2, 15–40. [405, 410, 411]

Dow, James and Gary Gorton (2008), "Noise traders." In *The New Palgrave Dictionary of Economics* (Steven N. Durlauf and Lawrence E. Blume, eds.), 54–57, Palgrave Macmillan, Basingstoke. [421]

Ely, Jeffrey C. and Marcin Pęski (2006), "Hierarchies of belief and interim rationalizability." *Theoretical Economics*, 1, 19–65. [411]

Frankel, David M., Stephen Morris, and Ady Pauzner (2003), "Equilibrium selection in global games with strategic complementarities." *Journal of Economic Theory*, 108, 1–44. [411, 416]

Klemperer, Paul (1998), "Auctions with almost common values: The 'wallet game' and its applications." *European Economic Review*, 42, 757–769. [420]

Kyle, Albert S. (1985), "Continuous auctions and insider trading." *Econometrica*, 53, 1315–1335. [421]

Mailath, G. and L. Samuelson (2006), *Repeated Games and Reputations: Long-Run Relationships*. Oxford University Press, New York. [421]

Mertens, Jean-François and Shmuel Zamir (1985), "Formulation of Bayesian analysis for games with incomplete information." *International Journal of Game Theory*, 14, 1–29. [409, 424]

Morris, Stephen and Hyun Song Shin (2003), "Global games: Theory and applications." In *Advances in Economics and Econometrics, Volume 1* (Mathias Dewatripont, Lars Peter Hansen, and Stephen J. Turnvosky, eds.), 56–114, Cambridge University Press, Cambridge. [408]

Oury, Marion and Olivier Tercieux (2012), "Continuous implementation." *Econometrica*, 80, 1605–1637. [427]

Pearce, David G. (1984), "Rationalizable strategic behavior and the problem of perfection." *Econometrica*, 52, 1029–1050. [411]

Penta, Antonio (2012a), "Higher order uncertainty and information: Static and dynamic games." *Econometrica*, 80, 631–660. [410, 417, 426]

Penta, Antonio (2012b), "Robust dynamic mechanism design." Unpublished paper, University of Wisconsin. [426]

Rubinstein, A. (1991), "The electronic mail game: Strategic behavior under 'almost common knowledge'." *American Economic Review*, 79, 385–391. [408]

Weinstein, Jonathan and Muhamet Yildiz (2007), "A structure theorem for rationalizability with application to robust predictions of refinements." *Econometrica*, 75, 365–400. [405, 407, 408, 410, 413, 416, 422, 426]

Weinstein, Jonathan and Muhamet Yildiz (2011), "Sensitivity of equilibrium behavior to higher-order beliefs in nice games." *Games and Economic Behavior*, 72, 288–300. [426]

Weinstein, Jonathan and Muhamet Yildiz (2012), "A structure theorem for rationalizability in infinite horizon-games." Unpublished paper. [426]

Submitted 2010-9-8. Final version accepted 2012-5-16. Available online 2012-5-16.