

HIERARCHIES OF BELIEF AND INTERIM RATIONALIZABILITY

JEFFREY C. ELY AND MARCIN PEŃSKI

ABSTRACT. In games with incomplete information, conventional hierarchies of belief are incomplete as descriptions of the players' information for the purposes of determining a player's behavior. We show by example that this is true for a variety of solution concepts. We then investigate what is essential about a player's information to identify rationalizable behavior in any game. We do this by constructing the universal type space for rationalizability and characterizing the types in terms of their beliefs. Infinite hierarchies of beliefs *over conditional beliefs*, what we call Δ -hierarchies, are what turn out to matter. We show that any two types in any two type spaces have the same rationalizable sets in all games if and only if they have the same Δ -hierarchies.

1. INTRODUCTION

1.1. **Example.** Consider the following two player game of incomplete information. There are two states of the world $\Omega = \{-1, +1\}$. Each player i has three actions $A_i = \{a_i, b_i, c_i\}$, and a payoff u_i which depends on the actions chosen by each player and the state of the world. The payoffs are summarized in the following table.

			a_2	b_2	c_2				a_2	b_2	c_2
a_1	1,1	-10,-10	-10,0			a_1	-10,-10	1,1	-10,0		
b_1	-10,-10	1,1	-10,0			b_1	1,1	-10,-10	-10,0		
c_1	0,-10	0,-10	0,0			c_1	0,-10	0,-10	0,0		
$\omega = +1$						$\omega = -1$					

Focusing only on the subsets $\{a_i, b_i\}$, we have a common interest game in which the players wish to choose the same action in the positive state and the opposite action in the negative state. Failing to coordinate is costly, and the action c_i is a “safe” alternative when, conditional on the state, i is uncertain of his opponent's action.

To complete the description of the game, we need to specify the players' information about the state, the players' information about one another's information, and so on. One way to

Financial support from NSF grant #9985462 and the Alfred P. Sloan foundation is gratefully acknowledged.

model these *hierarchies of belief* is given by the following type space: Each player i has two possible types, $T_i = \{-1, +1\}$, and there is a common prior $\mu \in \Delta(T_1 \times T_2 \times \Omega)$ given by

$$\mu(t_i, t_{-i}, \omega) = \begin{cases} \frac{1}{4} & \text{if } \omega = t_i \cdot t_{-i} \\ 0 & \text{otherwise} \end{cases}$$

In this type space, each player assigns each state equal *ex ante* probability, but would learn the state with certainty, were he to learn his opponent's type. In particular, if the two players' types have opposite sign, the state is certainly $\omega = -1$ and if the types have the same sign, then the state is certainly $\omega = +1$.

In the game with this information structure it is possible in Bayesian Nash equilibrium for the players to achieve perfect coordination where types $t_i = 1$ play a_i and types $t_i = -1$ play b_i . Symmetrically, there is another equilibrium where $t_i = 1$ play b_i and $t_i = -1$ play a_i . Obviously it is also an equilibrium for both to play c_i independent of type.

1.2. Type Spaces and Hierarchies. A type space is a convenient device for specifying in a parsimonious way the infinite string of data (hierarchies of belief) necessary to close the model. This was the view of Harsanyi (1967-68) and it is the standard practice in economic models with incomplete information. We can see from our example how hierarchies of beliefs are embedded in the type space. Each player assigns equal probability to each state of the world. These are the first-order beliefs. Since player i holds this first-order belief regardless of whether his type is -1 or $+1$, and since player $-i$ assigns probability 1 to player i having one of these two types, it follows that each player is certain of the others' first-order beliefs. These are the second-order beliefs. The same reasoning implies that each player is certain of the other's second-order beliefs and so on. Indeed, in this type space it is always *common knowledge* that the two states are equally likely.

One potential concern with the use of a type space for modeling hierarchies of belief is the following. If hierarchies of belief are what really matter, then we must be assured that any hierarchy we might wish to model can be captured in some type space. This concern has been resolved by Mertens and Zamir (1985) and Brandenburger and Dekel (1993) who showed that when the set of states of the world Ω has some minimal structure, then any internally consistent ("coherent") hierarchy can be modeled using a type space. In fact, there exists a single *universal* type space $U(\Omega)$ which simultaneously captures them all: for every coherent hierarchy there is a type in the universal type space with that hierarchy.¹

¹On the other hand, when the set of states lacks the topological structure assumed by these authors, Heifetz and Samet (1999) showed that the type space framework may not be sufficiently general to model all coherent hierarchies.

There is another potential concern which has not received the same attention. The type space we used in our example is but one of many that would capture those specific hierarchies of belief. Indeed, any specification of the players' hierarchies can be equally well generated by many different type spaces. If hierarchies are what matter, and if type spaces are simply a convenient device used to model them, then we should be assured that the predictions we would get from a given hierarchy should not depend on the particular type space used to model it. However, the type space can matter, and this can be seen in our example.

Recall that in our example, each player has the same hierarchy of beliefs regardless of his type. It appears that there is a spurious duplication of types. So instead consider the simpler type space in which each player has exactly one type and this type knows the other player's type and assigns equal probability to the two states of the world. Formally, $T_i^* = \{*\}$, and there is a common prior μ^* given by $\mu^*(*, *, +1) = \mu^*(*, *, -1)$. This type space generates exactly the same hierarchies of belief as in our first example: common knowledge that the states are equally likely. However, when the game in our example is played over this type space, the unique Bayesian Nash equilibrium, the unique correlated equilibrium, indeed the unique rationalizable outcome, is for both players to play c_i .²

We cannot be assured that our predictions are invariant to the choice of the type space. This is not a showstopper however, it simply means that our premise was wrong: it is *not* (only) hierarchies that matter for (correlated) equilibrium and rationalizability. While the additional types in the original type space are duplicates in terms of their hierarchies, they are not redundant because they generate a payoff-relevant means of correlating behavior *with the state of the world*.³ What matters in an incomplete information game are the hierarchies *and* the information a player with a given hierarchy would obtain about the state of the world conditional on knowing the other player has a given hierarchy.

²It deserves emphasis that the issue we are pointing to here is distinct from the familiar one that adding redundant types to an information structure creates the possibility that the players can correlate their action choices and thus increases the set of equilibrium outcomes. That observation is equivalent to the statement that the set of correlated equilibria of a game is larger than the set of Nash equilibria of a game. To see that something different is happening in our example, note that the sets of correlated equilibria in the two games are distinct. Adding redundant types in order to generate correlation in play can never affect the set of correlated equilibria (see Brandenburger and Dekel (1987)). Indeed, it can never affect the set of rationalizable outcomes as it does here.

³A similar example and observation appears in Dekel, Fudenberg, and Morris (2003). They introduce a new version of rationalizability in which players can conjecture correlations between the opponents action and the state beyond those correlations that are explicitly modeled in the type space. Conventional hierarchies are sufficient to identify the sets that are rationalizable under this alternative definition. In the context of our example, all actions satisfy their definition for every type in every type space.

To make the point in another way, let us frame the discussion in the context of the Mertens and Zamir (1985) and Brandenburger and Dekel (1993) universal type space. In $U(\Omega)$, there is exactly one type for each possible coherent hierarchy. And by the definition of a type space, for each type there is a belief defined over the product of the types of the other player and the states of the world. This implies that for each pair of hierarchies (i.e. universal type) t_i and t_{-i} , there is only one possible belief t_i can have about states of the world conditional on knowing that the opponent has hierarchy t_{-i} (modulo measure zero variations in versions of conditional probability). Indeed, one can interpret the representation theorems of Mertens and Zamir (1985) and Brandenburger and Dekel (1993) as proving that if the players' information is completely summarized by their hierarchies, then the modeler has no freedom in specifying these conditional beliefs as they are in fact uniquely determined. On the other hand, we see from our example that even in very simple, standard type spaces, two types with the same hierarchies can hold different conditional beliefs, and that this can make a difference in outcomes. It is easily verified that the universal types corresponding to the hierarchies in our example behave like the types in the second type-space, not the first.⁴ Thus, the universal type space is not rich enough to model what could be modeled in an alternative type space.

This observation has some significance for the philosophical debate (see Aumann (1987), Brandenburger and Dekel (1993), and Gul (1998)) about whether or not the information structure in a game is common knowledge. The universal type space has been interpreted as precisely that information structure that can be assumed without loss of generality to be common knowledge. For example, Brandenburger and Dekel (1993) suggest that the universal type space realizes Aumann's hypothesis of a completely specified "state space." This is certainly true if, as in Brandenburger and Dekel (1993), one considers the information structure purely as a model of beliefs (about beliefs) about uncertain events. But if what is important is the the range of possible behaviors in a game and not just beliefs, then our example shows that there *is* a loss of generality in assuming that the universal type space

⁴Any type in any type space has a "representative" type in $U(\Omega)$. The representative is the unique type in $U(\Omega)$ with the same hierarchy. Thus, any type space can be mapped into $U(\Omega)$. When this mapping applied to the first type-space we considered, the two types of player i are collapsed to a single type t_i^* which assigns probability 1 to t_{-i}^* . Thus, the image in $U(\Omega)$ of our first type space is isomorphic to our second type space.

is commonly known. In particular, this assumption would imply that whenever the players commonly know that each state in the example is equally likely, they must play action c .⁵

1.3. Universal Types for Rationalizability. As argued above, from the point of view of rationalizability, conventional hierarchies of belief are incomplete as descriptions of a player’s information. The goal of this paper is to identify exactly what must be known about a player’s information to determine what will be rationalizable for that player in any game of incomplete information. To do this, we first construct a type space which is universal in the following sense. Any type in any type space can be interpreted as a *rule* which associates each game form with the set of actions that are rationalizable for that type. We construct a space of types \mathcal{R} by collecting every rule associated with any type in any type space. We show how to view \mathcal{R} as a proper type space and that for this type space the rationalizable rule associated with any type is exactly the rule used to define it. If what we care about when we consider a player’s “type” is how that player might conceivably behave in any game (i.e. what is rationalizable), then \mathcal{R} is universal in the sense that it contains every possible type.

To further emphasize the point made in the previous subsections, let us observe that \mathcal{R} is larger than the standard universal type space. To see this, revisit our example from the introduction. Each of the types t_i of player i in the first type space we examined has a rationalizable rule. (In fact it is easily shown that they have the same rule.) Let us try to find a type in $U(\Omega)$ that gives rise to the same rule. The first thing to observe is that any such type must have the same hierarchy of beliefs as t_i . Indeed, for any two types with distinct hierarchies there is some game in which they have distinct rationalizable sets.⁶ But as we mentioned previously, there is only one type in $U(\Omega)$ with this hierarchy and this type cannot have the same rule because it has a different rationalizable set in the game from our example.

Next, we characterize the types in \mathcal{R} in terms of their beliefs. We show how to interpret any type t_i in any type space as an infinite hierarchy of beliefs *about conditional beliefs*. Here

⁵In Section 6 we present an example that makes an even stronger point in this regard. There we describe a game with an action that could not be played in any Bayesian Nash (or correlated) equilibrium by any type when it is assumed that the universal type space is common knowledge. Nevertheless, this action could be played in a very simple, completely standard type space where there is common knowledge of rationality and common knowledge that the players’ beliefs are coherent.

⁶This is a consequence of the more general Theorem 2 below. However, this particular implication can be shown directly by considering games (similar to those in Morris (2002)) in which players are asked to name their own hierarchies and bet on the announcement of the other. Two types with different hierarchies will disagree about some n -th order belief of the opponent and the state of the world, and there is some bet that would separate them.

the first order belief is the probability distribution over all the possible conditional beliefs the player could have about the state of the world. This probability distribution is derived for t_i as follows. First, determine for each type of the opponent what would be the conditional belief $\rho(t_i, t_{-i}) \in \Delta(\Omega)$ of t_i if the opponent's type were known. Then, the probability of any set $X \subset \Delta\Omega$ of possible conditional beliefs is the probability t_i assigns to the set of types t_{-i} for which $\rho(t_i, t_{-i})$ belongs to X . Once we have derived first-order beliefs of this form for every type, we can in the usual way derive the second-order beliefs: the probability any type t_i assigns to the events consisting of first-order beliefs of the opponent and the conditional beliefs of t_i . Higher-order beliefs are defined analogously. Let us refer to the resulting infinite hierarchies as Δ -hierarchies. In a completely standard fashion, we may construct the universal space $U(\Delta\Omega)$ of all Δ -hierarchies.

We show that the mapping which associates each type in \mathcal{R} with its Δ -hierarchy is a bijection, $\mathcal{R} \leftrightarrow U(\Delta\Omega)$. The implication is that within this particular type space, \mathcal{R} , to determine what is rationalizable for a player in any game, it is necessary and sufficient to identify that player's Δ -hierarchy. We then extend this characterization to all possible type spaces. We consider for any type space the natural mapping into \mathcal{R} that maps types to their associated rationalizable rule. This mapping preserves Δ -hierarchies. Thus, by the previous result, two types from any (possibly different) type spaces have the same rationalizable behavior across games if and only if they have the same Δ -hierarchies.

The rest goes as follows. Section 2 introduces notation and definitions used in the paper. Section 3 is central to the paper: it presents construction, analyzes properties and characterizes space \mathcal{R} . Sections 4 and 5 sketch some of the proofs. Section 6 comments and contains further examples.

2. PRELIMINARIES

2.1. Notation. If A is a topological space, we treat it as a measurable space with its Borel σ -field, denoted \mathcal{B}_A . ΔA is the space of Borel probability measures. If A is a Polish space, then ΔA endowed with the weak* topology is also Polish. For any $a \in A$ let $\delta(a) \in \Delta A$ be the Dirac measure concentrated on a point a .

For any measure $\mu \in \Delta A$ and integrable function $f : A \rightarrow \mathbf{R}$ we use $\mu[f]$ to denote the expectation of f with respect to μ . For any measure $\mu \in \Delta(A \times B)$, denote by $C_A\mu(\cdot) : B \rightarrow \Delta A$ a version of conditional probability over A given $b \in B$ (which exists whenever A is Polish. Our results do not depend on the choice of version). Similarly, for any measurable subset $B' \subseteq B$, we denote $C_A\mu(B') \in \Delta A$ as conditional probability given B' .

$\mathcal{K}A$ is the space of all non-empty closed subsets of A with the Hausdorff metric. If A is Polish, then so is $\mathcal{K}A$.

Given two measurable spaces, A, B and a measurable mapping $\phi : A \rightarrow B$ we can in a natural way define a mapping which transports probability measures $\Delta\phi : \Delta A \rightarrow \Delta B$, such that for any measure $\mu \in \Delta A$, any measurable subset $B' \subseteq B$, we have $\Delta\phi(\mu)(B') = \mu(\phi^{-1}(B'))$.

2.2. Type spaces. We consider games with two players. We take as given a Polish space of basic uncertainty Ω . A *Type space* over Ω , $T = (T_i, \mu_i)_{i=1,2}$ is a pair of measurable spaces T_i and two mappings $\mu_i : T_i \rightarrow \Delta(\Omega \times T_{-i})$. We say that a type space has *weakly measurable beliefs* if for any measurable function $f : \Omega \times T_{-i} \rightarrow \mathbf{R}$, sets

$$\{t_i : \mu[f] < 0\}$$

are measurable. We say that a type space has *strongly measurable beliefs*, if there exist jointly measurable functions $\rho_i : T_i \times T_{-i} \rightarrow \Delta\Omega$, such that

$$\rho_i(t_i, t_{-i}) = C_{\Omega}\mu(t_i)(t_{-i}).$$

Let $\mathcal{TW}(\Omega)$ be the collection of all type spaces over Ω with weakly measurable beliefs and $\mathcal{TS}(\Omega)$ be the collection of all type spaces over Ω with strongly measurable beliefs⁷.

A *type mapping* between two type spaces $T, T' \in \mathcal{TS}(\Omega)$, denoted $\phi : T \rightarrow T'$ is a pair of measurable mappings $\phi_i : T_i \rightarrow T'_i$.

2.3. Games. A *game form* (or simply game) over Ω is a tuple $G = (u_i, A_i)_i$, such that A_i are Polish spaces and $u_i : A_i \times A_{-i} \times \Omega \rightarrow \mathbf{R}$ are bounded measurable functions for both i . A game $G = (u_i, A_i)_i$ is *compact*, if u_i are continuous and A_i are compact. A game $G = (u_i, A_i)_i$ is *finite* if u_i are continuous and A_i are finite. We let \mathcal{G} denote the class of all compact games, however all of our results apply equally well if \mathcal{G} is taken to be (smaller) class of all finite games.

Sometimes it is useful to use *product games*: take any two games $G^1, G^2 \in \mathcal{G}$, $G^k = (A_i^k, u_i)$. We construct a product game $G = G_1 \times G_2 = (A_i, u_i)$, where the action sets in G are the products of the actions sets from the original games, $A_i = A_i^1 \times A_i^2$, and payoffs are given by

$$u_i((a_i^1, a_i^2), (a_{-i}^1, a_{-i}^2), \omega) = u_i^1(a_i^1, a_{-i}^1, \omega) + u_i^2(a_i^2, a_{-i}^2, \omega).$$

Note that $G \in \mathcal{G}$.

⁷Obviously any type space with strongly measurable beliefs has also weakly measurable beliefs, $\mathcal{TS}(\Omega) \subseteq \mathcal{TW}(\Omega)$. The connection in the other way is not clear. For any type space $T \in \mathcal{TW}(\Omega)$, standard theorems guarantee existence of conditional beliefs $\rho(t_i, t_{-i})$, which are measurable in t_{-i} for given t_i . We do not know, in general, whether we can choose conditional beliefs which are jointly measurable

This is in fact possible in some special cases. Suppose that sets of types T_i in type space $T \in \mathcal{TW}(\Omega)$ are separable and metrisable. Then, we think we may show that conditionals can be chosen in a jointly measurable way and $T \in \mathcal{TS}(\Omega)$.

2.4. Interim Rationalizability. Fix a type space $T \in \mathcal{TS}(\Omega)$, and a game $G = (A_i, u_i)$. An assessment is a pair of subsets $\alpha = (\alpha_1, \alpha_2)$ where $\alpha_i \subset T_i \times A_i$. Alternatively an assessment can be defined by the pair of correspondences $\alpha_i : T_i \rightrightarrows A_i$, with $\alpha_i(t_i) := \{a_i : (t_i, a_i) \in \alpha_i\}$. The image $\alpha_i(t_i)$ is interpreted as the set of actions that player i of type t_i could conceivably play.

A behavioral strategy for player i is a measurable function $\sigma_i : T_i \rightarrow \Delta A_i$. The expected payoff to type t_i of player i from choosing action a_i when the opponent's strategy is σ_{-i} is given by

$$\begin{aligned} U_i(a_i, \sigma_{-i}|t_i) &= \mu_i^T(t_i)[\sigma_{-i}(t_{-i})[u_i(a_i, \cdot, \cdot)]] \\ &= \int_{T_{-i} \times \Omega} \int_{A_{-i}} u_i(a_i, \cdot, \omega) d\sigma_{-i}(t_{-i}) d\mu_i^T(t_i) \end{aligned}$$

The strategy σ_i is a selection from the assessment α if for each i , $\sigma_i(t_i) \in \Delta \alpha_i(t_i)$ for all $t_i \in T_i$. Let $\Sigma_i(\alpha_i)$ be the set of all strategies for i that are selections from α .

For some results, it is convenient to use an alternative notation for payoffs and strategies. Given a payoff function $u_i : A \times \Omega \rightarrow \mathbf{R}$, we derive a new payoff function $\pi_i : A \times T \rightarrow \mathbf{R}$, defined directly in terms of the types as follows

$$\pi_i(a, t) = \rho_i(t_i, t_{-i})[u_i(a_i, a_{-i}, \cdot)].$$

A *conjecture* for player i is a probability measure $\sigma_{-i}^\Delta \in \Delta(T_{-i} \times A_{-i})$. The set of conjectures for a given type t_i is the set $\Sigma^\Delta(t_i)$ of σ_{-i}^Δ satisfying

$$\text{marg}_{T_{-i}} \sigma_{-i}^\Delta = \text{marg}_{T_{-i}} \mu_i^T(t_i).$$

For any behavior strategy σ_{-i} , and for any t_i , there is a conjecture $\sigma_{-i}^\Delta \in \Sigma^\Delta(t_i)$ such that

$$U_i(a_i, \sigma_{-i}|t_i) = \sigma_{-i}^\Delta[\pi_i(a_i, \cdot, t_i, \cdot)] := U_i(a_i, \sigma_{-i}^\Delta|t_i)$$

for every $a_i \in A_i$. Conversely, if σ_{-i}^Δ is a conjecture for t_i , then there is a behavior strategy σ_{-i} satisfying the same equalities. This is to show that we can either work with behavioral strategies or conjectures, whichever is most convenient.

An action a_i is an interim best-response for t_i against σ_{-i} if $U_i(a_i, \sigma_{-i}|t_i) \geq U_i(a'_i, \sigma_{-i}, |t_i)$ for all $a'_i \in A_{-i}$. Let $B_i(\sigma_{-i}|t_i)$ denote the set of all interim best-responses for t_i to σ_{-i} . Likewise $B_i^\Delta(\sigma_{-i}^\Delta|t_i)$ the set of all best-responses to the conjecture σ_{-i}^Δ . If X is some subset of strategies (or conjectures), then $B_i(X|t_i)$ (respectively $B_i^\Delta(X|t_i)$) is the set of all best-responses to elements of X .

An assessment α has the *best-response property* if every action attributed to player i is an interim best-reply to some selection from α_{-i} , i.e.,

$$\alpha_i \subset \{(t_i, a_i) : a_i \in B_i(\Sigma_{-i}(\alpha_{-i})|t_i)\}$$

If the above is satisfied with equality, then we say that α has the fixed-point property.

Proposition 1. *There exists a maximal (in the sense of set inclusion) assessment with the fixed-point property*

Proof. It is easy to verify that the union of assessments with the best-response property has the best-response property. Let R be the union of all assessments with the best-response property. Obviously R is the maximal set with the best-response property. We claim that R has the fixed-point property, in which case it will be the maximal such set. If R did not have the fixed-point property then there exists a type t_i and action a_i such that a_i is an interim best-reply to some selection from R_{-i} . But then we can add the pair (t_i, a_i) to R and obtain a larger assessment with the best-response property, a contradiction. \square

Definition 1. *Given a type space T , and a game G , the interim rationalizable correspondence is the maximal assessment with the fixed-point property, denoted $R^{G,T}$. We say that a_i is interim rationalizable for type t_i if $a_i \in R_i^{G,T}(t_i)$.*

We conclude this section by stating some important properties of rationalizable correspondences which are proved in Section 5. The rationalizable correspondence is non-empty, closed-valued, and measurable in a strong sense.⁸

Proposition 2. *For each type space $T \in \mathcal{TS}(\Omega)$, for each type $t_i \in T_i$, the set $R_i^{G,T}(t_i)$ of interim rationalizable actions is non-empty and closed. Thus, we can view $R_i^{G,T}$ as a function from T_i to \mathcal{KA}_i . This function is measurable: for every $\mathcal{B} \in \mathcal{B}_{\mathcal{KA}_i}$, the set*

$$\{t_i \in T_i : R_i^{G,T}(t_i) \in \mathcal{B}\}$$

is a measurable subset of T_i .

3. THE SPACE OF RATIONALIZABLE RULES

In this section we construct and characterize the universal type space of rationalizable rules.

3.1. Construction. We begin with the following sets for each i :

$$\mathcal{S}_i = \prod_{G \in \mathcal{G}} \mathcal{KA}_i^G$$

Any element r_i of \mathcal{S}_i can be viewed as a rule which assigns a (closed) subset of A_i^G to each $G \in \mathcal{G}$ - recall that \mathcal{KA}_i^G is a compact Polish space with the Hausdorff metric. The value of

⁸It is somewhat surprising that the correspondence need not be closed (i.e. upper hemi-continuous) even for compact games and Ω compact, indeed even for finite games when Ω is finite. This is shown by example in section 6.

a rule r_i on particular game $G \in \mathcal{G}$ is denoted by $r_i(G)$. We equip \mathcal{S}_i with the associated product topology and Borel σ -algebra.

A rule r_i is *rationalizable* if there exists a type space T and a type $t_i \in T_i$ such that $R_i^{G,T}(t_i) = r_i(G)$ for every $G \in \mathcal{G}$. We use the notation $R_i^{(\cdot),T} : T_i \rightarrow \mathcal{S}_i$ for the mapping which associates types in T with their corresponding rules. Let \mathcal{R}_i be the subset of \mathcal{S}_i consisting of all rationalizable rules, i.e.

$$\mathcal{R}_i = \bigcup_{T \in \mathcal{TS}(\Omega)} R_i^{(\cdot),T}(T_i).$$

The space \mathcal{R}_i inherits the topology and σ -algebra from \mathcal{S}_i . The derived σ -algebra is denoted $\mathcal{B}_{\mathcal{R}_i}$.

We pause here to record the following important fact used repeatedly later.

Lemma 1. *For any type space $T \in \mathcal{TS}(\Omega)$, the mapping $R_i^{(\cdot),T} : T_i \rightarrow \mathcal{R}_i$ is measurable.*

Proof. By the monotone class theorem and the choice of topology on \mathcal{R}_i , we need to check that for any finite number of games $G^1, \dots, G^k \in \mathcal{G}$ and open set $KA \subseteq \mathcal{KA}_i^{G^1} \times \dots \times \mathcal{KA}_i^{G^k}$, the set

$$\left(R_i^{(\cdot),T} \right)^{-1} \left(KA \times \prod_{G \in \mathcal{G}, G \neq G_i \text{ for } i=1, \dots, k} \mathcal{KA}_i^G \right) = \left(R_i^{G^1,T} \times \dots \times R_i^{G^k,T} \right)^{-1} (KA)$$

is measurable. Consider the product game $G = G^1 \times \dots \times G^k = (A_i^G, u_i^G)$. Observe that KA is an open subset $KA \subseteq \mathcal{KA}_i^G$ and $R_i^{G^1,T} \times \dots \times R_i^{G^k,T} = R_i^{G,T}$. Now by proposition 2, the set $\left(R_i^{G,T} \right)^{-1} (KA)$ is measurable in T_i . \square

3.2. \mathcal{R} as a Type Space. Our goal is to treat \mathcal{R} as a type space by assigning beliefs to each rule $r_i \in \mathcal{R}_i$. We do this by transporting in the natural way the beliefs of some type t_i (in some existing type space) from which r_i was derived. There are two purposes of this exercise. First, it will allow us to check whether the construction is internally consistent. Having specified beliefs of each type in \mathcal{R} we can solve for the rationalizable actions *within the type space* \mathcal{R} of any type r_i for any game G and determine the rule $R_i^{(G),\mathcal{R}}(r_i)$ generated by r_i . Then we can check whether this rule $R_i^{(G),\mathcal{R}}(r_i)$ indeed coincides with r_i . Next, once we have verified the internal consistency, we can proceed to characterize rationalizable rules in terms of their beliefs and higher-order beliefs. The latter will be the focus of the remainder of the paper.

There is an important subtlety involved in transforming \mathcal{R} in an internally consistent way into a type space. In this subsection we will revisit the example from the introduction to illustrate the issues involved and to motivate the development to follow. Consider the direct approach to defining the beliefs of a rule $r_i \in \mathcal{R}_i$. Since r_i is a rationalizable rule, there

is a type t_i in a type space T which generates r_i . Recall that $R_{-i}^{(\cdot),T} : T_{-i} \rightarrow \mathcal{R}_{-i}$ is the measurable mapping which carries types from T_{-i} into the space \mathcal{R}_{-i} . To define the beliefs of type $r_i = R_i^{(\cdot),T}(t_i)$ over the types of the opponent, we similarly transport the belief of t_i :

$$\text{marg}_{\mathcal{R}_{-i}} \mu_i^{\mathcal{R}}(r_i) = \left(\Delta R_{-i}^{(\cdot),T} \right) \left(\text{marg}_{T_{-i}} \mu_i^T(t_i) \right) \quad (3.1)$$

This is not enough however, as we must define beliefs of r_i over types of the opponent *and* states of the world. To make these beliefs consistent with those of the original type t_i , the only choice is to put:

$$C_{\Omega} \mu_i^{\mathcal{R}}(r_i)(r_{-i}) = C_{\Omega} \mu_i^T(t_i) \left(\left(R_{-i}^{(\cdot),T} \right)^{-1}(r_{-i}) \right) \quad (3.2)$$

i.e. conditional on a rule r_{-i} of the opponent, r_i holds the same beliefs over Ω as t_i would have conditional on the set $(R_{-i}^{(\cdot),T})^{-1}(r_{-i})$ of types in T_{-i} which generate rule r_{-i} .

Unfortunately, with beliefs translated in this way, some important information from the original type space T is lost, and the internal consistency of the construction fails. To see this, let us trace through these mappings as applied to the type space from the introduction. First of all, as we mentioned previously, both types of player i generate the same rationalizable rule, so $R_i^{(\cdot),T}(+1) = R_i^{(\cdot),T}(-1) := r_i$ for $i = 1, 2$. Then by (3.1), type r_i must assign probability 1 to the single type r_{-i} (she knows the rule of her opponent), and by (3.2), type r_i must assign equal probability to each state of the world (conditional on r_{-i}). In other words, independent of the opponent's information, r_i thinks that each state is equally likely. But such a type is isomorphic to the type $*$ from the second type space T^* in the example which we previously showed has a different rationalizable rule. In the terminology we use below, this way of mapping beliefs does not preserve rationalizable sets and this means that the rule of type r_i is not equal to r_i , i.e. the model is not internally consistent.

The right way to transform \mathcal{R} into a type space is therefore somewhat more involved. Our development will involve two steps. First, we need to identify what information contained in the beliefs from the original type space T is important for rationalizability in that type space. As our discussion in the introduction suggests, what are important are beliefs about *conditional* beliefs. To develop a framework for such beliefs, we will consider a new class of type spaces where the space of basic uncertainty is $\Delta\Omega$, the space of (conditional) beliefs about Ω . A type in such a type space has beliefs about the opponent's type and $\Delta\Omega$. In the usual fashion, these can be unfolded into beliefs about $\Delta\Omega$, beliefs about the opponent's beliefs about $\Delta\Omega$ and so on. We will show how to treat the space of rationalizable rules as a type space $\mathcal{R}(\Delta)$ over $\Delta\Omega$. In a certain sense, this is the right way to think about the space of rules.

Nevertheless, in order to determine what is rationalizable for a type in a game (whose payoffs are defined over Ω) we must be able to calculate expected payoffs and so we must interpret the beliefs as probabilities over the opponent's type and Ω . Having seen the construction of the space $\mathcal{R}(\Delta)$, it will be readily apparent how to transform $\mathcal{R}(\Delta)$ back into a type space over Ω in such a way that the final type space, $\mathcal{R}(\Omega)$, is internally consistent.

3.3. The Type Space $\mathcal{R}(\Delta)$. Let us take as our space of basic uncertainty the (Polish) space $\Delta\Omega$. Exactly as with Ω itself, we can define a type space over $\Delta\Omega$. The beliefs of a type t_i in such a type space T are probabilities over $\Delta(\Delta\Omega \times T_{-i})$. We interpret these as joint probabilities over the types of the opponent and conditional beliefs about Ω . We will consider the class of all type spaces over $\Delta\Omega$ with weakly measurable beliefs, denoted $\mathcal{TW}(\Delta\Omega)$.

There is a natural way in which any type space $T = (T_i, \mu_i)$ over Ω can be transformed into a type space $T^\Delta = (T_i^\Delta, \mu_i^\Delta) \in \mathcal{TW}(\Delta\Omega)$. Let $T_i^\Delta = T_i$ and for any $t_i \in T_i^\Delta$ we define $\mu_i^\Delta(t_i) \in \Delta(\Delta\Omega \times T_{-i})$ to be the unique probability measure satisfying the following two conditions:

- (1) beliefs about opponent types are unchanged,

$$\text{marg}_{T_{-i}^\Delta} \mu_i^\Delta(t_i) = \text{marg}_{T_{-i}} \mu_i(t_i),$$

- (2) conditional beliefs of $t_i \in T_i^\Delta$ about $\Delta\Omega$ given type t_{-i} are a point mass on the conditional belief t_i in T_i given type t_{-i}

$$C_{\Delta\Omega} \mu_i^\Delta(t_i)(t_{-i}) = \delta(C_{\Omega} \mu_i(t_i)(t_{-i})),$$

here $\delta(\cdot)$ denotes Dirac delta measure.

The logic behind this definition is the following. The translated beliefs $\mu_i^\Delta(t_i)$ capture exactly the joint probability over opponent's types and the resulting conditional beliefs as embodied in $\mu_i(t_i)$. In the appendix we show that this defines a weakly measurable belief mapping.

Lemma 2. *Suppose that $T \in \mathcal{TS}(\Omega)$. Then $T^\Delta \in \mathcal{TW}(\Delta\Omega)$*

Now, we may define the beliefs of a rule $r_i \in \mathcal{R}_i$ generated by some type t_i by transporting the transformed beliefs $\mu_i^\Delta(t_i)$. Find some type space T and type $t_i \in T_i$, such that $R_i^{(\cdot), T}(t_i) = r_i$ and put

$$\mu_i^R(r_i) = \Delta \left(\text{id}_{\Delta\Omega} \times R_{-i}^{(\cdot), T} \right) (\mu_i^\Delta(t_i)).$$

As defined, the beliefs of type-rules r_i are potentially dependent on the choice of type space and type $t_i \in T_i$. Since there are potentially many type spaces with different types

having the same rationalizable rule, it is important to verify that all such types will generate the same beliefs:

Proposition 3. *For any two type spaces $T, T' \in \mathcal{TS}(\Omega)$, any two types $t_i \in T_i, t'_i \in T'_i$, if both types generate the same rules,*

$$R_i^{(\cdot),T}(t_i) = R_i^{(\cdot),T'}(t'_i),$$

then they also generate the same beliefs about the conditional beliefs and the opponent's rules,

$$\Delta \left(\text{id}_{\Delta\Omega} \times R_{-i}^{(\cdot),T} \right) \left(\mu_i^\Delta(t_i) \right) = \Delta \left(\text{id}_{\Delta\Omega} \times R_{-i}^{(\cdot),T'} \right) \left(\mu_i^\Delta(t'_i) \right).$$

To see how this treatment avoids the problem with the more direct approach, let us return to the example from the introduction. Recall that r_i is the rule generated by both types of player i . The previous proposition states that w.l.o.g the beliefs of r_i can be derived from type $+1$. The first step is to derive the belief μ_i^Δ :

$$\mu_i^\Delta(+1, \delta_{\{\omega=+1\}}) = \mu_i^\Delta(-1, \delta_{\{\omega=-1\}}) = 1/2.$$

Next, we apply the mapping $\Delta \left(\text{id}_{\Delta\Omega} \times R_{-i}^{(\cdot),T} \right)$ to obtain

$$\mu_i^{\mathcal{R}}(r_i)(r_{-i}, \delta_{\{\omega=+1\}}) = \mu_i^{\mathcal{R}}(r_i)(r_{-i}, \delta_{\{\omega=-1\}}) = 1/2.$$

Just as before, the two types of the opponent generate the same rule and so will be mapped to a single element r_{-i} and so r_i must assign probability 1 to r_{-i} , i.e. she knows the rule of her opponent. However, now type r_i assigns equal probability to the events “the opponent has rule r_{-i} and I will know for sure that the state is $+1$ ” and “the opponent has rule r_{-i} and I will know for sure that the state is -1 .” This is exactly the joint probability over (\mathcal{R} -equivalence classes of) types of the opponent and conditional beliefs held by the original type from which r_i was derived.

Having defined the beliefs, the final element is to show that the belief mapping is weakly measurable, so that $\mathcal{R}(\Delta)$ is a well-defined type space. The proof of the following proposition is deferred to the next section (4.2).

Proposition 4. *The belief mapping $\mu_i^{\mathcal{R}} : \mathcal{R}_i \rightarrow \Delta(\Delta\Omega \times \mathcal{R}_{-i})$ is weakly measurable. Thus, $\mathcal{R} \in \mathcal{TW}(\Delta\Omega)$*

3.4. The Type Space $\mathcal{R}(\Omega)$. By viewing \mathcal{R} as a type space over $\Delta\Omega$, we have identified rationalizable rules with their beliefs about the types of the opponent and conditional probabilities of events in Ω . Recall that Proposition 3 states that for type spaces over Ω , these beliefs precisely determine a type's rationalizable rule. Our goal is to show that the mapping which associates an arbitrary type t_i with its rule $r_i = R_i^{(\cdot),T}(t_i) \in \mathcal{R}(\Delta)$ preserves (the essential structure of) these beliefs and therefore that as a type $R_i^{(\cdot),T}(t_i)$ has the

same rationalizable rule as t_i , namely r_i . This will show that the construction is internally consistent.

However we cannot make this argument working with $\mathcal{R}(\Delta)$, as it is not a type space over Ω . First, to be able to discuss rationalizability, we must find the right model for $\mathcal{R}(\Delta)$ within the class of type spaces over Ω . The structure of beliefs in $\mathcal{R}(\Delta)$ suggests the answer. In $\mathcal{R}(\Delta)$, given a rule for the opponent, a type r_i has a (possibly non-degenerate) probability distribution over conditional beliefs about Ω . On the other hand, within any type space over Ω , any type can have only one conditional belief about Ω for each possible type of the opponent. Thus, to replicate the structure of beliefs from $\mathcal{R}(\Delta)$ within a type space over Ω , we must include a sufficiently rich set of types. Indeed, our set of types will be of the form $\mathcal{R}_i \times \Delta\Omega$ with the idea that type (r_i, τ) will be the type whose rationalizable rule will be r_i and conditional on which player $-i$ will have belief τ over Ω .

Precisely, we define a type space over Ω , $\mathcal{R}(\Omega) = (\mathcal{R} \times \Delta\Omega, \mu_i^{\mathcal{R}(\Omega)})$ in the following way. Let $(\mathcal{R} \times \Delta\Omega)_i = \mathcal{R}_i \times \Delta\Omega$ be the space of types. For every type $(r_i, \tau_i) \in \mathcal{R}_i \times \Delta\Omega$ we define the beliefs

$$\mu_i^{\mathcal{R}(\Omega)}(r_i, \tau_i) \in \Delta(\Omega \times (\mathcal{R} \times \Delta\Omega)_{-i})$$

as the unique measure satisfying the following two conditions:

- (1) marginal beliefs about $\Delta\Omega \times \mathcal{R}_{-i}$ are obtained from $\mathcal{R}(\Delta)$,

$$\text{marg}_{\Delta\Omega \times \mathcal{R}_{-i}} \mu_i^{\mathcal{R}(\Omega)}(r_i, \tau_i) = \mu_i^{\mathcal{R}}(r_i),$$

- (2) given any $(r_{-i}, \tau_{-i}) \in (\mathcal{R} \times \Delta\Omega)_{-i}$ conditional beliefs about Ω are equal to τ_{-i} ,

$$C\mu_i^{\mathcal{R}(\Omega)}(r_i, \tau_i)(r_{-i}, \tau_{-i}) = \tau_{-i}$$

Obviously the conditional probabilities are measurable so that these beliefs properly define a type space with strongly measurable beliefs, $\mathcal{R}(\Omega) \in \mathcal{TS}(\Omega)$.

3.5. Internal Consistency of $\mathcal{R}(\Omega)$. For any type space $T \in \mathcal{TS}(\Omega)$ and an arbitrary $\tau \in \Delta\Omega$, consider the type-mapping $\phi^{T,\tau} : T \rightarrow \mathcal{R} \times \Delta\Omega$ defined by

$$\phi_i^{T,\tau}(t_i) = (R_i^{(\cdot),T}(t_i), \tau). \quad (3.3)$$

The measurability of $\phi^{T,\tau}$ follows from Lemma 1.

Definition 2. We say that type mapping $\phi : T \rightarrow T'$, $T, T' \in \mathcal{TS}(\Omega)$ preserves rationalizable sets, if for all games $G \in \mathcal{G}$, for all types $t_i \in T_i$

$$R_i^{G,T}(t_i) = R_i^{G,T'}(\phi_i(t_i)).$$

We will present some sufficient conditions for a mapping to preserve rationalizable sets and show that $\phi^{T,\tau}$ satisfies these conditions. These conditions are described in terms of the way in which beliefs are implicitly transformed by the mapping.

Suppose that we have two type spaces $T = (T_i, \mu_i), T' = (T'_i, \mu'_i)$ over the same space of basic uncertainty X (for our purposes, X will be either Ω or $\Delta\Omega$.) Consider a type mapping $\phi : T \rightarrow T'$, such that for every t_i , there is a measurable mapping $\phi_{t_i} : T_{-i} \rightarrow T'_{-i}$, such that, for any measurable subset $S' \subseteq T'_{-i}$

$$\text{marg}_{T'_{-i}} \mu'(\phi(t_i))(S') = \text{marg}_{T_{-i}} \mu(t_i)(\phi_{t_i}^{-1}(S')). \quad (3.4)$$

The mappings ϕ_{t_i} are referred to as the *dual* mappings. We can interpret the dual mapping ϕ_{t_i} as describing how player i type t_i “thinks” that types of the opponent are transported. We say that a type mapping is *exact* iff $\phi_{t_i} = \phi$ for every t_i (or types t_{-i} are transported in exactly the same way as type t_i ”thinks”). We say that a type mapping ϕ is *exact with respect to beliefs* iff for every $t_i \in T_i$ and every $t_{-i} \in T_{-i}$,

$$\mu'(\phi_{t_i}(t_{-i})) = \mu'(\phi(t_{-i})) \quad (3.5)$$

(or *beliefs* of types t_{-i} are transported in exactly the same way as type t_i ”thinks” about transporting beliefs). Thus, an exact mapping is also exact with respect to beliefs. We say that a type mapping *preserves beliefs* iff for any type $t_i \in T_i$ and for almost any type $t'_{-i} \in T'_{-i}$

$$C_X \mu(t_i)(\phi_{t_i}^{-1}(t'_{-i})) = C_X \mu'(\phi(t_i))(t'_{-i}). \quad (3.6)$$

We say that a type mapping *preserves conditional beliefs* iff for any type $t_i \in T_i$, for almost any type $t_{-i} \in T_{-i}$

$$C_X \mu(t_i)(t_{-i}) = C_X \mu'(\phi(t_i))(\phi_{t_i}(t_{-i})). \quad (3.7)$$

Thus, if a mapping preserves conditional beliefs, then it also preserves beliefs. However, the converse does not hold necessarily. Consider the type mapping informally described in the introduction which associated types in the example with their counterparts in the universal type space $U(\Omega)$. It can easily be checked that this mapping preserves beliefs but does not preserve conditional beliefs. Indeed the mappings considered by Mertens and Zamir (1985) and Heifetz and Samet (1999) which associate types with their hierarchies are equivalent to our exact, belief-preserving type-mappings. They do not preserve conditional beliefs. The following lemma shows that the latter is a sufficient condition for a mapping to preserve rationalizable sets.

Lemma 3. *Suppose that for type spaces $T, T' \in \mathcal{TS}(\Omega)$ there is a type mapping $\phi : T \rightarrow T'$, which is exact with respect to beliefs and preserves conditional beliefs. Then it preserves rationalizable sets.*

Proof. Fix a game G . The rationalizable correspondence for G in type space T is $R_i^{G,T}$. Consider the following assessment for typespace T'

$$\alpha'_i = \hat{\phi}(R_i^{G,T}) \cup \bigcup_{t'_{-i} \in T'_{-i}} \hat{\phi}_{t'_{-i}}(R_i^{G,T})$$

where $\hat{\phi} = \phi \times id_A$. We will show that α'_i has the best-response property.

Let us write $S'_i = \phi(T_i) \cup \bigcup_{t'_{-i}} \phi_{t'_{-i}}(T_i)$. Note that for any $t'_i \in S'_i$,

$$\alpha'_i(t'_i) = \bigcup_{t_i \in \phi^{-1}(t'_i)} R_i^{G,T}(t_i) \cup \bigcup_{t'_{-i}} \bigcup_{t_i \in \phi_{t'_{-i}}^{-1}(t'_i)} R_i^{G,T}(t_i) \quad (3.8)$$

and for any $t'_i \notin S'_i$, $\alpha'_i(t'_i) = \emptyset$.

Pick $t_i \in T_i$. Let $t'_i = \phi(t_i)$, $a_i \in R_i^{G,T}(t_i)$, and $\sigma_{-i}^{\Delta,T} \in \Sigma_{-i}^{\Delta,T}(R_{-i}^{G,T}|t_i)$ with $a_i \in B_i^T(\sigma_{-i}^{\Delta,T})$. We construct a conjecture for t'_i as follows:

$$\sigma_{-i}^{\Delta,T'} = \sigma_{-i}^{\Delta,T} \circ \hat{\phi}_i^{-1}$$

We claim that for any $z_i \in A_i$,

$$U_i^{T'}(z_i, \sigma_{-i}^{\Delta,T'}) = U_i^T(z_i, \sigma_{-i}^{\Delta,T})$$

To show this, we first use the fact that ϕ preserves conditional beliefs to establish that the type-dependent payoff function π_i is preserved under ϕ . For any action profile a , and type profile $\hat{t} \in T$,

$$\begin{aligned} \pi_i^T(a, \hat{t}) &= (C_\Omega \mu_i^T(\hat{t}_i) (\hat{t}_{-i})) [u_i(a, \cdot)] \\ &= \left(C_\Omega \mu_i^{T'}(\phi(\hat{t}_i)) (\phi_{t_i}(\hat{t}_{-i})) \right) [u_i(a, \cdot)] \\ &= \pi_i^{T'}(a, \phi(\hat{t}_i), \phi_{t_i}(\hat{t}_{-i})) \end{aligned}$$

Next, it follows that

$$\begin{aligned} U_i^T(z_i, \sigma_{-i}^{\Delta,T}) &= \sigma_{-i}^{\Delta,T} [\pi_i^T(z_i, a_{-i}, \hat{t})] \\ &= \sigma_{-i}^{\Delta,T} [\pi_i^{T'}(z_i, a_{-i}, \phi(\hat{t}_i), \phi_{t_i}(\hat{t}_{-i}))] \\ &= \sigma_{-i}^{\Delta,T'} [\pi_i^{T'}(z_i, a_{-i}, \cdot, \cdot)] \\ &= U_i^{T'}(z_i, \sigma_{-i}^{\Delta,T'}). \end{aligned}$$

The third equality holds because by the construction of $\sigma_{-i}^{\Delta,T}$, for any measurable subset $C \subset \mathbf{R}$,

$$\sigma_{-i}^{\Delta,T'}(\{(t'_{-i}, a_{-i}) : \pi_i^{T'}(z_i, a_{-i}, t'_i, t'_{-i}) \in C\}) = \sigma_{-i}^{\Delta,T}(\{(t_{-i}, a_{-i}) : \pi_i^T(z_i, a_{-i}, t'_i, \phi_{t_i}(t_{-i})) \in C\}).$$

This establishes our claim.

Next, note that $\sigma_{-i}^{\Delta, T'} \in \Sigma_{-i}^{\Delta, T'}(\alpha'_{-i}|t'_i)$. In particular, because $\phi_{t_i}(R_{-i}^{G, T}) \subset \alpha'_{-i}$, we have $\sigma_{-i}^{\Delta, T'}(\alpha'_{-i}) \geq \sigma_{-i}^{\Delta, T'}(R_{-i}^{G, T}) = 1$. We have therefore shown that $a_i \in B_i^{T'}(\alpha'_{-i}|t'_i)$. since a_i was arbitrary, we conclude $R_i^{G, T}(t_i) \subset B_i^{T'}(\alpha'_{-i}|t'_i)$.

Next consider $t''_i = \phi_{t_{-i}}(t_i)$ for some t_{-i} . Because ϕ is exact with respect to beliefs, $\mu_i^{T'}(t'_i) = \mu_i^{T'}(t''_i)$. It follows that

$$R_i^{G, T}(t_i) \subset B_i^{T'}(\alpha'_{-i}|t'_i) = B_i^{T'}(\alpha'_{-i}|t''_i)$$

It now follows from (3.8) that α' has the best-response property.

Now turn to the rationalizable correspondence $R_i^{G, T'}$ on T' . Construct an assessment for T as follows.

$$\alpha_i(t_i) = \bigcup_{t_{-i}} R_i^{G, T'}(\phi_{t_{-i}}(t_i))$$

We will show that α has the best-response property.

Pick $t_i \in T_i$. Because ϕ is exact with respect to beliefs, for any t_{-i} , $\phi_{t_{-i}}(t_i)$ and $\phi(t_i)$ have the same beliefs in T_{-i} , and therefore $R_i^{G, T'}(\phi_{t_{-i}}(t_i)) = R_i^{G, T'}(\phi(t_i))$. Thus, $\alpha_i(t_i) = R_i^{G, T'}(\phi(t_i))$. Let $a_i \in R_i^{G, T'}(\phi(t_i))$ so that there is a conjecture $\sigma_{-i}^{\Delta, T'} \in \Sigma_{-i}^{\Delta, T'}(R_{-i}^{G, T'}|\phi(t_i))$ such that $a_i \in B_i^{T'}(\sigma_{-i}^{\Delta, T'})$. We derive a conjecture for t_i in a few steps.

First, consider the behavior strategy σ'_{-i} defined by $\sigma'_{-i}(t'_{-i}) = C_{A_{-i}}\sigma_{-i}^{\Delta, T'}(t'_{-i})$. We can select a measurable version of $C_{A_{-i}}\sigma_{-i}^{\Delta, T'}$ so that $\sigma'_{-i}(t'_{-i}) \in \Delta R_{-i}^{G, T'}(t'_{-i})$ for each t'_{-i} . Now there corresponds a behavior strategy for T , namely $\sigma_{-i} = \sigma'_{-i} \circ \phi_{t_i}$. Construct the conjecture $\sigma_{-i}^{\Delta, T}$ from $\mu_i^T(t_i)$ and σ_{-i} . It can easily be checked that this construction yields

$$\sigma_{-i}^{\Delta, T'} = \sigma_{-i}^{\Delta, T} \circ \hat{\phi}_i^{-1}$$

and since $\alpha_{-i} = \hat{\phi}_i^{-1}(R_{-i}^{G, T'})$,

$$\sigma_{-i}^{\Delta, T}(\alpha_{-i}) = [\sigma_{-i}^{\Delta, T} \circ \hat{\phi}_i^{-1}](R_{-i}^{G, T'}) = \sigma_{-i}^{\Delta, T'}(R_{-i}^{G, T'}) = 1$$

hence $\sigma_{-i}^{\Delta, T} \in \Sigma_{-i}^{\Delta, T}(\alpha_{-i}|t_i)$.

Furthermore, the claim from the first half of the proof applies and we can conclude $a_i \in B_i^T(\sigma_{-i}^{\Delta, T})$ so that $a_i \in B_i^T(\alpha_{-i}|t_i)$. As a_i was selected arbitrarily, we have shown that α has the best-response property.

To summarize, we have shown that there is an assessment on T' with the best-response property such that $R_i^{G, T}(t_i) \subset \alpha'_i(\phi(t_i))$. Since $R_i^{G, T'}$ includes any assessment with the best-response property, it follows that

$$R_i^{G, T}(t_i) \subset R_i^{G, T'}(\phi(t_i))$$

Likewise, in the opposite direction we showed

$$R_i^{G, T'}(\phi(t_i)) \subset R_i^{G, T}(t_i)$$

and this concludes the proof. \square

It is easy to check that by the definition of beliefs in the type space $\mathcal{R}(\Delta)$, for any type space $T \in \mathcal{TS}(\Omega)$, the type mapping $R^{(\cdot),T} : T^\Delta \rightarrow \mathcal{R}(\Delta)$ is in fact exact and preserves beliefs. We argue in the proposition below that, as a consequence, the type mapping $\phi^{T,\tau} : T \rightarrow \mathcal{R}(\Omega)$ is exact with respect to beliefs and preserves conditional beliefs. Then lemma 3 allows us to say that $\phi^{T,\tau}$ preserves rationalizable sets.

Proposition 5. *For any type space $T \in \mathcal{TS}(\Omega)$ and $\tau \in \Delta\Omega$, the type mapping $\phi^{T,\tau} : T \rightarrow \mathcal{R}(\Omega)$ is exact with respect to beliefs and preserves conditional beliefs. Hence, it preserves rationalizable sets. Moreover, for any rule $r_i \in \mathcal{R}_i$, any $\tau_i \in \Delta\Omega$*

$$r_i = R_i^{(\cdot),\mathcal{R}(\Omega)}(r_i, \tau_i).$$

Proof. Let $T = (T_i, \mu_i)$. We define the dual mapping for type $t_i \in T_i$ by

$$\phi_{t_i}^{T,\tau}(t_{-i}) = \left(R_{-i}^{(\cdot),T}(t_{-i}), C_\Omega \mu_i(t_i)(t_{-i}) \right).$$

By Lemma 1, $\phi_{t_i}^{T,\tau}$ is measurable. We check that for any measurable set $B \subset (\mathcal{R} \times \Delta\Omega)_{-i}$

$$\begin{aligned} \text{marg}_{\mathcal{R}_{-i} \times \Delta\Omega} \mu_i^{\mathcal{R}(\Omega)} \left(\phi_{t_i}^{T,\tau}(t_i) \right) (B) &= \text{marg}_{\mathcal{R}_{-i} \times \Delta\Omega} \mu_i^{\mathcal{R}(\Omega)} \left(R_{-i}^{(\cdot),T}(t_i), \tau \right) (B) \\ &= \mu_i^{\mathcal{R}} \left(R_{-i}^{(\cdot),T}(t_i) \right) (B) \\ &= \mu_i^{T^\Delta}(t_i) \left(\left\{ (t_{-i}, \tau) : \left(R_{-i}^{(\cdot),T}(t_{-i}), \tau \right) \in B \right\} \right) \\ &= \text{marg}_{T_{-i}} \mu_i(t_i) \left(\left\{ t_{-i} : \left(R_{-i}^{(\cdot),T}(t_{-i}), C_\Omega \mu_i(t_i)(t_{-i}) \right) \in B \right\} \right) \\ &= \text{marg}_{T_{-i}} \mu_i(t_i) \left(\phi_{t_i}^{-1}(B') \right) \end{aligned}$$

(the third equality comes from our observation in the text that the mapping $R^{(\cdot),T} : T^\Delta \rightarrow \mathcal{R}(\Delta)$ is exact and preserves beliefs; the fourth from the definition of beliefs on the space T^Δ). Therefore, the type mapping $\phi_{t_i}^{T,\tau}$ satisfies equation (3.4) and so is a valid dual mapping. Verification that $\phi^{T,\tau}$ preserves conditional beliefs becomes straightforward:

$$\begin{aligned} C_\Omega \mu_i^{\mathcal{R}(\Omega)} \left(\phi_{t_i}^{T,\tau}(t_i) \right) \left(\phi_{t_i}^{T,\tau}(t_{-i}) \right) &= C_\Omega \mu_i^{\mathcal{R}(\Omega)} \left(\phi_{t_i}^{T,\tau}(t_i) \right) \left(R_{-i}^{(\cdot),T}(t_{-i}), C_\Omega \mu_i(t_i)(t_{-i}) \right) \\ &= C_\Omega \mu_i(t_i)(t_{-i}). \end{aligned}$$

In order to check exactness with respect to beliefs, note that for any $r_i \in \mathcal{R}_i$, any $\tau_i, \tau'_i \in \Delta\Omega$, $\mu_i^{\mathcal{R}(\Omega)}(r_i, \tau_i) = \mu_i^{\mathcal{R}(\Omega)}(r_i, \tau'_i)$ so that for any $t_i \in T_i$, $t_{-i} \in T_{-i}$

$$\begin{aligned} \mu_i^{\mathcal{R}(\Omega)} \left(\phi_{t_i}^{T,\tau}(t_{-i}) \right) &= \mu_i^{\mathcal{R}(\Omega)} \left(R_{-i}^{(\cdot),T}(t_{-i}), \tau \right) \\ &= \mu_i^{\mathcal{R}(\Omega)} \left(R_{-i}^{(\cdot),T}(t_{-i}), C_\Omega \mu_i(t_i)(t_{-i}) \right) = \mu_i^{\mathcal{R}(\Omega)} \left(\phi_{t_i}^{T,\tau}(t_{-i}) \right). \end{aligned}$$

Take now any rationalizable rule $r_i \in \mathcal{R}_i$ and some $\tau_i \in \Delta\Omega$. There is type in a type space $t_i \in T_i$, such that $r_i = R_i^{(\cdot),T}(t_i)$. Since the type mapping $\phi^{T,\tau}$ preserves rationalizable sets, we have

$$r_i = R_i^{(\cdot),T}(t_i) = R_i^{(\cdot),\mathcal{R}(\Omega)}(r_i, \tau) = R_i^{(\cdot),\mathcal{R}(\Omega)}(r_i, \tau_i).$$

The last equality comes from the fact that beliefs of types in $\mathcal{R}(\Omega)$, hence by proposition 3 also their rationalizable sets, depend only on the r_i -coordinate. \square

The construction of the space $\mathcal{R}(\Omega)$ assures that types $(r_i, \tau_i), (r'_i, \tau'_i) \in \mathcal{R}_i(\Omega)$ have the same beliefs if they agree on the first coordinate, $r_i = r'_i$. The second part of the last proposition strengthens the implication into equivalence: since types $(r_i, \tau_i), (r'_i, \tau'_i) \in \mathcal{R}_i(\Omega)$ must have different rationalizable behavior whenever $r_i \neq r'_i$, it follows that they must have also different beliefs. Since beliefs of $(r_i, \tau_i) \in \mathcal{R}(\Omega)$ are generated from beliefs of $r_i \in \mathcal{R}(\Delta)$, it follows that two different rules $r_i \neq r'_i$ must have different beliefs in type space $\mathcal{R}(\Delta)$. In particular, the implication of proposition 3 may be strengthened to equivalence: two types generate the same rules if and only if they generate the same beliefs about conditional beliefs and rules of the opponent.

3.6. Δ -hierarchies. Our construction of the space of rationalizable rules has one advantage and one fault. The advantage is that it is direct and relatively simple. The fault is that its construction reveals little about the internal structure of the space. The goal of this section is to characterize rules in terms of their beliefs. We do this by first presenting an alternative type space which is defined directly in terms of players' beliefs. Then we show that this space is actually equivalent to $\mathcal{R}(\Delta)$.

The idea behind this alternative construction is very simple. Recall that $\mathcal{R}(\Delta)$ is a type space over $\Delta\Omega$. In completely standard fashion we may derive from the beliefs of any type its first-order beliefs about $\Delta\Omega$, second-order beliefs about the opponent's first-order beliefs cross $\Delta\Omega$, etc. In other words, types in $\mathcal{R}(\Delta)$ can be interpreted as infinite hierarchies of beliefs over the space $\Delta\Omega$. We refer to these as Δ -hierarchies. We present below a version of the Mertens and Zamir (1985) theorem on the existence of a universal type space $U(\Delta\Omega)$ consisting of *all* such hierarchies.

Theorem 1 (Mertens and Zamir (1985), Brandenburger and Dekel (1993), Battigalli and Siniscalchi (1999)). *Let X be a Polish space and $\mathcal{TW}(X)$ the class of all type spaces over X with weakly measurable beliefs. There exists a universal type space $U(X) \in \mathcal{TW}(X)$ such that for any type space $T \in \mathcal{TW}(X)$, there is a unique exact and beliefs preserving type*

mapping $u_i^T : T_i \rightarrow U_i(X)$ ⁹. Moreover, $U_i(X) \simeq \Delta(X \times U_{-i}(X))$ for both players i (in the sense of homeomorphism).

The proof is in the appendix. It follows that of Mertens, Sorin, and Zamir (1994). The latter applies to continuous type spaces, but the proof is easily adapted to cover measurable type spaces as in our case. When we apply the theorem to the Polish space $\Delta\Omega$, the mapping $u_i^{T\Delta} : T_i^\Delta \rightarrow U(\Delta\Omega)$ gives the Δ -hierarchy of type $t_i \in T_i$ for any type space $T \in \mathcal{TS}(\Omega)$.

We will show that $\mathcal{R}(\Delta)$ and $U(\Delta\Omega)$ are isomorphic with respect to exact and belief-preserving type mappings. First, note that just as we transformed $\mathcal{R}(\Delta)$ into a type space over Ω , we may define a type space $L(\Omega)$ where the set of types for i is $U_i(\Delta\Omega) \times \Delta\Omega$. and the belief mapping $\mu_i^{L(\Omega)}$ is defined by

$$\begin{aligned} \text{marg}_{\Delta\Omega \times U_{-i}^\Delta} \mu_i^{L(\Omega)}(u_i, \tau_i) &= \mu_i^{U(\Delta\Omega)}(u_i) \\ C_\Omega \mu_i^{L(\Omega)}(u_i, \tau_i)(u_{-i}, \tau_{-i}) &= \tau_{-i} \end{aligned}$$

for any $(u_i, \tau_i) \in L_i^\Omega$, and $(u_{-i}, \tau_{-i}) \in L_{-i}^\Omega$. In particular, notice that beliefs (hence also rationalizable sets) of any type $(u_i, \tau_i) \in L_i^\Omega$ depend only on its u_i -coordinate.

We show how to map any type space $T \in \mathcal{TS}(\Omega)$ into $L(\Omega)$ in a way that preserves rationalizable sets. First, fix some $\tau \in \Delta\Omega$ and let $\text{in}_i^\tau : U_i(\Delta\Omega) \rightarrow L_i(\Omega)$ be the inclusion mapping $\text{in}_i^\tau(u_i) = (u_i, \tau)$. We may define a type mapping $l_i^{T,\tau} : T \rightarrow L(\Omega)$ as the following composition

$$l^{T,\tau} = \text{in}_i^\tau \circ u_i^{T\Delta}. \quad (3.9)$$

Such a mapping preserves rationalizable sets, as the following lemma together with lemma 3 show:

Lemma 4. *For any type space $T \in \mathcal{T}(\Omega)$, a type mapping $l^{T,\tau} : T \rightarrow L(\Omega)$ is exact with respect to beliefs and preserves conditional beliefs.*

Proof. The proof here is completely analogous to the first half of the proof in proposition 5 □

Before we can show that spaces $\mathcal{R}(\Delta)$ and $U(\Delta\Omega)$ are equivalent (or, in other words, the spaces $\mathcal{R}(\Omega)$ and $L(\Omega)$) are equivalent, we need one more result. Define the projection $\text{proj}_T : T \times \Delta\Omega \rightarrow T$ with $\text{proj}(r_i, \tau_i) = r_i$.

Lemma 5. *Let $S, T \in \mathcal{TS}(\Delta\Omega)$.*

⁹In the language of category theory, universal type space is a *terminal* object within the category of type spaces $\mathcal{TW}(X)$ connected with type morphisms which are exact and preserve beliefs.

(1) For any exact, belief preserving type mapping $\phi^\Delta : S \rightarrow T$, mapping

$$\phi = \text{in}^{T,\tau} \circ \phi^\Delta \circ \text{proj}_S$$

is an exact with respect to beliefs, conditional beliefs preserving type mapping $\phi : S \times \Delta\Omega \rightarrow T \times \Delta\Omega$.

(2) Suppose T has the following property: $t_i \neq t'_i \implies \mu_i(t_i) \neq \mu_i(t'_i)$. Then for any exact with respect to beliefs, conditional beliefs preserving type mapping $\phi : S \times \Delta\Omega \rightarrow T \times \Delta\Omega$, mapping

$$\phi^\Delta = \text{proj}_T \circ \phi \circ \text{in}^{S,\tau}$$

is an exact, belief preserving type mapping $\phi^\Delta : S \rightarrow T$.

Notice that both spaces, $\mathcal{R}(\Delta)$ and $U(\Delta\Omega)$ have the property assumed in the second part of lemma: $U(\Delta\Omega)$ by the previous theorem and $\mathcal{R}(\Delta)$ by the remark following proposition 5.

Proof. We begin with the first part of lemma: ϕ^Δ is exact and preserves beliefs. For any $(s_i, \tau_i) \in S_i \times \Delta\Omega$, define $\phi_{(s_i, \tau_i)} : S_{-i} \times \Delta\Omega \rightarrow T_{-i} \times \Delta\Omega$ with

$$\phi_{(s_i, \tau_i)}(s_{-i}, \tau_{-i}) = (\phi^\Delta(s_{-i}), \tau_{-i}).$$

We check that (3.4) holds: for any subset $B' \subseteq T_i \times \Delta\Omega$

$$\begin{aligned} \text{marg}_{T_{-i} \times \Delta\Omega} \mu^{T \times \Delta\Omega}(\phi(s_i, \tau_i))(B') &= \text{marg}_{T_{-i} \times \Delta\Omega} \mu^{T \times \Delta\Omega}(\phi^\Delta(s_i), \tau)(B') \\ &= \mu^T(\phi^\Delta(s_i))((t_i, \tau_{-i}) \in B') \\ &= \mu^{S_i}(s_i)((\phi^\Delta(s_{-i}), \tau_{-i}) \in B') \\ &= \mu^{S_i}(s_i)((\phi^\Delta(s_{-i}), \tau_{-i}) \in B') \\ &= \text{marg}_{S_{-i} \times \Delta\Omega} \mu^{S \times \Delta\Omega}(s_i, \tau_i)(\phi_{(s_i, \tau_i)}^{-1}(B')) \end{aligned}$$

(where the third equality comes from exactness and belief preserving of ϕ^Δ). It is immediate to verify that ϕ is exact with respect to beliefs and preserves conditional beliefs.

We move to the second part of the lemma: Since ϕ preserves conditional beliefs and is exact with respect to beliefs, there is for any $(s_i, \tau_i) \in S \times \Delta\Omega$ a measurable mapping $\phi_{(s_i, \tau_i)} : S_{-i} \times \Delta\Omega \rightarrow T_{-i} \times \Delta\Omega$, such that equations (3.4), (3.5), (3.7) hold. Due to the fact that ϕ is exact with respect to beliefs, it must be that if $\phi_{-i}(s_{-i}, \tau_{-i}) = (t_{-i}, \tau_{-i})$, then for every $(s_i, \tau_i) \in S_i \times \Delta\Omega$, there is $\tau'_i \in \Delta\Omega$, such that $\phi_{(s_i, \tau_i)}(s_{-i}, \tau_{-i}) = (t_i, \tau'_i)$. Indeed, types with different t -coordinate have different beliefs. Therefore,

$$\text{proj}_T \circ \phi_{(s_i, \tau_i)} = \text{proj}_T \circ \phi$$

and ϕ^Δ is exact. Since ϕ preserves conditional beliefs, it must be that $\phi(s_i, \tau_i) = (\phi_i^\Delta(s_i), \tau_i)$. This implies the second equality in the following: for any $s_i \in S_i$, for any $t_{-i} \in T_{-i}$

$$\begin{aligned} C_{\Delta\Omega} \mu_i^S(s_i) \left((\phi_i^\Delta)^{-1}(t_{-i}) \right) &= C_{\Delta\Omega} \operatorname{marg}_{S_{-i} \times \Delta\Omega} \mu_i^{S \times \Delta\Omega}(s_i, \tau) \{ (\phi_i)^{-1}(\{t_{-i}\} \times \Delta\Omega) \} \\ &= C_{\Delta\Omega} \operatorname{marg}_{T_{-i} \times \Delta\Omega} \mu_i^{T \times \Delta\Omega}(\phi_i(s_i, \tau)) (\{t_{-i}\} \times \Delta\Omega) \\ &= C_{\Delta\Omega} \mu_i^T(\phi_i^\Delta(s_i))(t_{-i}), \end{aligned}$$

and ϕ^Δ preserves beliefs. \square

We can finally prove the theorem

Theorem 2. *There are unique exact and beliefs preserving type mappings $\iota : U(\Delta\Omega) \rightarrow \mathcal{R}(\Delta)$ and $\iota^{-1} : \mathcal{R}(\Delta) \rightarrow U(\Delta\Omega)$. Either mapping is inverse of the other: $\iota^{-1} \circ \iota = \operatorname{id}_{U(\Delta\Omega)}$ and $\iota \circ \iota^{-1} = \operatorname{id}_{\mathcal{R}(\Delta)}$.*

Proof. Existence and uniqueness of type mapping ι^{-1} is assured by theorem 1. In order to show existence of exact, beliefs-preserving mapping from $U(\Delta\Omega)$ to $\mathcal{R}(\Delta)$, note first that by the first half of proposition 5, there is exact with respect to beliefs and conditional beliefs preserving type mapping $\phi^{L(\Omega), \tau} : L(\Omega) \rightarrow \mathcal{R}(\Omega)$. The second part of lemma 5 then guarantees existence of exact and beliefs preserving mapping $(\phi^{L(\Omega), \tau})^\Delta : U(\Delta\Omega) \rightarrow \mathcal{R}(\Delta)$.

Suppose now that we have two different exact and belief preserving mappings $\iota_1, \iota_2 : U(\Delta\Omega) \rightarrow \mathcal{R}(\Delta)$. There is $u_i \in U_i(\Delta\Omega)$, such that $\iota_1(u_i) \neq \iota_2(u_i)$. By the first part of lemma 5 there are then two type mappings $\phi_1, \phi_2 : L(\Omega) \rightarrow \mathcal{R}(\Omega)$ which preserve conditional beliefs and $\phi_1(u_i, \tau) = (i_1(u_i), \tau) \neq (i_2(u_i), \tau) = \phi_2(u_i, \tau)$. By lemma 3, $\phi_1(u_i, \tau)$ and $\phi_2(u_i, \tau)$ must have the same rationalizable rules. But this is a contradiction because proposition 5 shows that the rationalizable rule for $\phi_1(u_i, \tau)$ is $\iota_1(u_i)$ while the rationalizable rule for $\phi_2(u_i, \tau)$ is $\iota_2(u_i)$.


The equality $\iota^{-1} \circ \iota = \operatorname{id}_{U(\Delta\Omega)}$ comes from the uniqueness of exact and belief preserving mapping from $U(\Delta\Omega)$ to itself. The second equality $\iota \circ \iota^{-1} = \operatorname{id}_{\mathcal{R}(\Delta)}$ is a consequence of the fact that $\iota \circ \iota^{-1}$ would generate an exact with respect to beliefs and conditional belief preserving mapping from $\mathcal{R}(\Omega)$ to itself (the first part of lemma 5 guarantees that). Such a mapping has to preserve rules, so it has to preserve r -coordinates of types in $\mathcal{R}(\Omega)$ This implies that $\iota \circ \iota^{-1} = \operatorname{id}_{\mathcal{R}(\Delta)}$ \square

This leads directly to the following corollary which is the main result of the paper.

Corollary 1. *Any two types from any two type spaces have the same rationalizable rules if and only if they have the same Δ -hierarchies, i.e.*

$$R_i^{(\cdot), T}(t_i) = R_i^{(\cdot), S}(s_i) \iff u_i^{T, \Delta}(t_i) = u_i^{S, \Delta}(s_i)$$

Proof. Consider the following diagram, where $\tilde{\iota} = \text{in}^{\mathcal{R},\tau} \circ \iota \circ \text{proj}_U$ (as in the first part of Lemma 5.)



The route $T \rightarrow T^\Delta \rightarrow U(\Delta\Omega) \rightarrow L(\Omega)$ is the mapping $l^{T,\tau}$ (see 3.9). By Lemma 4, this mapping is exact with respect to beliefs and preserves conditional beliefs. By the first part of Lemma 5, $\tilde{\iota}$ has the same properties. Since these properties are preserved under composition, the route $T \rightarrow T^\Delta \rightarrow U(\Delta\Omega) \rightarrow L(\Omega) \rightarrow \mathcal{R}(\Omega)$ is exact with respect to beliefs and preserves conditional beliefs, hence it preserves rationalizable sets. Thus, it defines the same mapping as the direct route $\phi^{T,\tau}$. This, together with Theorem 2 establishes that the diagram is commutative. Thus, we can determine the rationalizable rules of types t_i and s_i by tracing the indirect route to $\mathcal{R}(\Omega)$. Because ι is a bijection, the result follows immediately. \square

4. BELIEF MAPPING

We have two goals in this section. First we want to show that belief mapping on the space of rules does not depend on the choice of representative type and type space - hence it is well-defined (proposition 3). We show that types with the same rule have the same beliefs about $\Delta\Omega$ and rules of the opponent and types with different beliefs have different rules. To this end we use special characterization theorem, proof of which, among others, contains a construction of games which have different rationalizable sets for different beliefs. Second, we attempt to show that the belief mapping is measurable (4).

4.1. Characterization of rules through beliefs. In the previous section we defined beliefs of type-rules over conditional beliefs and type-rules of the opponent, $\Delta\Omega \times \mathcal{R}_{-i}$. It is often useful to consider only restricted version of these beliefs: only about conditional beliefs and rationalizable sets of the opponent in some particular game $G = (A_i^G, u_i^G)$, i.e. beliefs about $\Delta\Omega \times \mathcal{K}A_{-i}^G$. Precisely, for any type we may define beliefs $\Delta R_i^{G,T}(t_i) \in \Delta(\Delta\Omega \times \mathcal{K}A_{-i}^G)$ with

formula:

$$\Delta R^{G,T}(t_i) := \Delta \left(\text{id}_{\Delta\Omega} \times R_{-i}^{G,T} \right) \left(\mu_i^\Delta(t_i) \right).$$

In other words, in order to determine beliefs $\Delta R^{G,T}(t_i)$, we need to take first type $t_i \in T_i$, find her beliefs $\mu_i^{T^\Delta}(t_i) \in \Delta(\Delta\Omega \times T_{-i})$, compute rationalizable sets of the correspondence of the opponent types in game G and transport naturally beliefs $\mu_i^{T^\Delta}(t_i)$ into beliefs $\Delta \left(\text{id}_{\Delta\Omega} \times R_{-i}^{G,T} \right) \left(\mu_i^\Delta(t_i) \right)$.

In particular, if proposition 3 is true and we may properly define beliefs for every rule, beliefs $\Delta R_i^{G,T}(t_i)$ are equal to marginal beliefs of the type-rule. Precisely, take any game $G \in \mathcal{G}$, any type $t_i \in T_i$ and observe that, if proposition 3 holds, then

$$\Delta R_i^{G,T}(t_i) = \underset{\Delta\Omega \times \mathcal{K}A_{-i}^G}{\text{marg}} \mu_i^{\mathcal{R}(\Delta)} \left(R_i^{(\cdot),T}(t_i) \right) = \underset{\Delta\Omega \times \mathcal{K}A_{-i}^G}{\text{marg}} \mu_i^{\mathcal{R}(\Delta)}(r_i)$$

for rationalizable rule $r_i = R_i^{(\cdot),T}(t_i)$. This comes from the definition of the space of rules: its elements are product of rationalizable sets in different games from \mathcal{G} .

The important fact is:

Theorem 3. *For any two type spaces $T, T' \in \mathcal{T}(\Omega)$, any two types, $t_i \in T_i$ and $t'_i \in T'_i$, if for every game $G \in \mathcal{G}$, rationalizable sets of both types are equal,*

$$R_i^{G,T}(t_i) = R_i^{G,T'}(t'_i),$$

then for every game $G \in \mathcal{G}$, beliefs of both types about $\Delta\Omega$ and rationalizable sets of the opponent in G are equal,

$$\Delta R_i^{G,T}(t_i) = \Delta R_i^{G,T'}(t'_i).$$

In other words, if rationalizable sets of both types are equal for any game G then types' beliefs about rationalizable sets of the opponent and conditional beliefs about state of the world are equal for any game G . (Note the order of the quantifiers!). We show first how we can use the theorem to prove proposition 3:

Proof of Proposition 3. The proposition is a consequence of the theorem, but with one complication. The theorem implies that for any game $G \in \mathcal{G}$

$$\begin{aligned} \underset{\Delta\Omega \times \mathcal{K}A_i^G}{\text{marg}} \Delta \left(\text{id}_{\Delta\Omega} \times R_{-i}^{(\cdot),T} \right) \left(\mu_i^\Delta(t_i) \right) &= \Delta R_i^{G,T}(t_i) = \Delta R_i^{G,T'}(t'_i) \\ &= \underset{\Delta\Omega \times \mathcal{K}A_i^G}{\text{marg}} \Delta \left(\text{id}_{\Delta\Omega} \times R_{-i}^{(\cdot),T} \right) \left(\mu_i^\Delta(t'_i) \right), \end{aligned}$$

or that marginal beliefs about conditional beliefs and opponent rationalizable sets in game G are the same. This statement holds for any G . However, it is not enough for us, since we need to argue that joint beliefs about opponent's rules (all of their rationalizable sets) are actually the same.

By Kolmogorov consistency theorem, it is enough to show that for any finite number of games $G_1, \dots, G_k \in \mathcal{G}$ we have

$$\text{marg}_{\Delta\Omega \times \mathcal{K}A_i^{G_1} \dots \times \mathcal{K}A_i^{G_k}} \Delta \left(\text{id}_{\Delta\Omega} \times R_{-i}^{(\cdot), T} \right) (\mu_i^\Delta (t_i)) = \text{marg}_{\Delta\Omega \times \mathcal{K}A_i^{G_1} \dots \times \mathcal{K}A_i^{G_k}} \Delta \left(\text{id}_{\Delta\Omega} \times R_{-i}^{(\cdot), T} \right) (\mu_i^\Delta (t'_i)).$$

The last equality is equivalent to

$$\Delta \left(\text{id}_{\Delta\Omega} \times R_{-i}^{G_1, T} \times \dots \times R_{-i}^{G_k, T} \right) (\mu_i^\Delta (t_i)) = \Delta \left(\text{id}_{\Delta\Omega} \times R_{-i}^{G_1, T} \times \dots \times R_{-i}^{G_k, T} \right) (\mu_i^\Delta (t'_i)).$$

To show this is however not so difficult: Observe that product of rationalizable sets in games G_1, \dots, G_k is equal to the set of rationalizable actions in the product game $G = G_1 \times \dots \times G_k$, where product game is defined in the section 2. This means that

$$\begin{aligned} & \Delta \left(\text{id}_{\Delta\Omega} \times R_{-i}^{G_1, T} \times \dots \times R_{-i}^{G_k, T} \right) (\mu_i^\Delta (t_i)) \\ &= \Delta R_i^{G, T} (t_i) = \Delta R_i^{G, T'} (t'_i) \\ &= \Delta \left(\text{id}_{\Delta\Omega} \times R_{-i}^{G_1, T} \times \dots \times R_{-i}^{G_k, T} \right) (\mu_i^\Delta (t'_i)). \end{aligned}$$

which ends the proof. \square

We prove now the theorem. Suppose that there are two types $t_i \in T_i, t'_i \in T'_i$ which have different beliefs about $\Delta\Omega$ and rationalizable sets in at least one game $G \in \mathcal{G}$. We show that it means that there is potentially another game $G' \in \mathcal{G}$, in which each of two types has different rationalizable sets. The difficulty is that G' is not necessarily equal to G - it is possible that two types have the same rationalizable sets of actions in same game, despite differences in beliefs.

Before we continue, we need a technical result. For a fixed game $G = (A_j, u_j) \in \mathcal{G}$, define the following set of continuous bounded functions $f : \Delta\Omega \times \mathcal{K}A_{-i} \rightarrow [0, \infty)$

$$\mathcal{F}^G = \left\{ \begin{array}{l} f(\tau, K) = \max_{k=1, \dots, N_1} \sup_{a_1, \dots, a_{N_2} \in K} \tau [\psi(k, a_1, \dots, a_{N_2}, \omega)] : \\ \text{for some natural } N_1, N_2 \text{ and} \\ \text{continuous bounded function } \psi : \{1, \dots, N_1\} \times A_{-i}^{N_2} \times \Omega \rightarrow [0, \infty). \end{array} \right\}.$$

On the first coordinate $\tau \in \Delta\Omega$, functions $f \in \mathcal{F}^G$ are ‘‘piecewise linear’’ and convex. On the second coordinate $K \in \mathcal{K}A_{-i}$, they are set-increasing: for any two sets $K \subseteq K', K, K' \in \mathcal{K}A_{-i}$, $f(\tau, K) \leq f(\tau, K')$. Next, define the set of differences of functions from \mathcal{F}^G

$$\mathcal{L}^G = \{f - g : f, g \in \mathcal{F}^G\} \subseteq C(\Delta\Omega \times \mathcal{K}A_{-i}).$$

We have the following lemma:

Lemma 6. *For any game G the collection of sets $\{\mu : \mu[f] < 0\} \subseteq \Delta(\Delta\Omega \times \mathcal{K}A_{-i})$ for $f \in \mathcal{L}^G$ generates the weak* topology on $\Delta(\Delta\Omega \times \mathcal{K}A_{-i})$. In particular, for any $\mu, \mu' \in \Delta(\Delta\Omega \times \mathcal{K}A_{-i})$, $\mu \neq \mu'$ there is function $f^G \in \mathcal{F}^G$, such that*

$$\mu[f^G(\tau, A)] \neq \mu'[f^G(\tau, A)].$$

We leave the proof of lemma for the appendix.

Proof of Theorem 3: Suppose that $\Delta R^{G^*, T} \mu(t_i) \neq \Delta R^{G^*, T'} \mu'(t'_i)$ for some game $G^* = (A_j^*, u_j^*)$. The lemma says that there are natural numbers N_1, N_2 and continuous bounded function $\psi : \{1, \dots, N_1\} \times A_{-i}^{N_2} \times \Omega \rightarrow [0, \infty)$ such that for $f : \Delta\Omega \times \mathcal{K}A_{-i} \rightarrow \mathbf{R}$ defined by

$$f(\tau, K) = \max_{k=1, \dots, N_1} \sup_{a_1, \dots, a_{N_2} \in K} \tau[\psi(k, a_1, \dots, a_{N_2}, \omega)],$$

we have either

$$\Delta R^{G^*, T} \mu(t_i)[f] < \Delta R^{G^*, T'} \mu'(t'_i)[f]$$

or

$$\Delta R^{G^*, T} \mu(t_i)[f] > \Delta R^{G^*, T'} \mu'(t'_i)[f].$$

Suppose w.l.o.g. that the first strict inequality holds. Find $\lambda > 0$, such that

$$\Delta R^{G^*, T} \mu(t_i)[\lambda f - 1] < 0 < \Delta R^{G^*, T'} \mu'(t'_i)[\lambda f - 1]$$

We will prove the theorem by constructing a game $G = (A_j, u_j)$, such that $R^{G, T}(t_i) \neq R^{G, T'}(t'_i)$.

First, find a game $G_0 = (A_j^0, u_j^0)$, such that $A_{-i}^0 = \{1, \dots, N_1\}$ and all actions of player $-i$ are rationalizable for all types of player $-i$, i.e. for any $t_{-i} \in T_{-i}$ any $t'_{-i} \in T'_{-i}$, $R_{-i}^{G_0, T}(t_{-i}) = R_{-i}^{G_0, T'}(t'_{-i}) = A_{-i}^0$ (such a game always exists). Denote $Z = \{0, 1\}$ and define sets of actions in game G as

$$\begin{aligned} A_i &= A_i^0 \times (A_i^*)^{N_2} \times Z, \\ A_{-i} &= A_{-i}^0 \times (A_{-i}^*)^{N_2}. \end{aligned}$$

Payoffs of player $-i$ are given by

$$\begin{aligned} &u_{-i}((a_{-i}^0, a_{-i,1}^*, \dots, a_{-i, N_2}^*), (a_i^0, a_{i,1}^*, \dots, a_{i, N_2}^*, z), \omega) \\ &= u_{-i}^0(a_{-i}^0, a_i^0, \omega) + \sum_{k=1}^{N_2} u_{-i}^*(a_{-i, k}^*, a_{i, k}^*, \omega) \end{aligned}$$

(in particular they do not depend on z) and payoffs of player i are given by

$$\begin{aligned} & u_i \left((a_i^0, a_{i,1}^*, \dots, a_{i,N_2}^*, z), (a_{-i}^0, a_{-i,1}^*, \dots, a_{-i,N_2}^*), \omega \right) \\ &= u_i^0(a_i^0, a_{-i}^0, \omega) + \sum_{k=1}^{N_2} u_i^*(a_{i,k}^*, a_{-i,k}^*, \omega) + z \left[\lambda \psi(a_{-i}^0, a_{-i,1}^*, \dots, a_{-i,N_2}^*, \omega) - 1 \right]. \end{aligned}$$

We show that the rationalizable sets for types t_i and t'_i are different in G . First observe, that due to the product structure of the game G , for any type space S , for any type $s_{-i} \in S_{-i}$,

$$R_{-i}^{G,S}(s_{-i}) = A_{-i}^0 \times \left[R_{-i}^{G^*,S}(s_{-i}) \right]^{N_2}$$

and for any $s_i \in S_i$,

$$R_i^{G,S}(s_i) = R_i^{G_0}(s_i) \times \left[R_i^{G^*,S}(s_i) \right]^{N_2} \times Z_i(s_i),$$

where $Z_i(s_i) \subseteq Z$.

In type space T , consider the (pure) behavioral strategy for player $-i$ which as type s_{-i} selects $a_{-i}^0 \in \{1, \dots, N_1\}$ and $(a_{-i,1}^*, \dots, a_{-i,N_2}^*)$ from $R_{-i}^{G^*,S}(s_{-i})$ to maximize the expression $\rho_i(t_i)(s_{-i})[\psi(a_{-i}^0, a_{-i,1}^*, \dots, a_{-i,N_2}^*, \omega)]$. By the measurable maximum theorem, this defines a measurable selection from $R_{-i}^{G^*,T}$.¹⁰ Call this strategy σ_{-i} . We can define the analogous strategy σ'_{-i} for type space T' where type t'_i replaces t_i in the definition.

We calculate the payoff to type t'_i of player i from playing $z = 1$ against σ'_{-i} :

$$\begin{aligned} & \mu'_i(t'_i) \left[\lambda \cdot \max_{k=1, \dots, N_1} \sup_{a_1, \dots, a_{N_2} \in K} \rho_i^{T'}(t'_i)(s'_{-i}) [\psi(k, a_1, \dots, a_{N_2}, \omega)] - 1 \right] \\ &= \Delta R^{G^*,T'} \mu'(t'_i) [\lambda f - 1] > 0. \end{aligned}$$

Thus, $1 \in Z_i(t'_i)$. On the other hand, the strategy σ_{-i} clearly maximizes, among all rationalizable strategies for player $-i$ in type space T , the payoff that type t_i could receive from playing $z = 1$ and

$$\begin{aligned} & \mu_i(t_i) \left[\lambda \cdot \max_{k=1, \dots, N_1} \sup_{a_1, \dots, a_{N_2} \in K} \rho_i^T(t_i)(s_{-i}) [\psi(k, a_1, \dots, a_{N_2}, \omega)] - 1 \right] \\ &= \Delta R^{G^*,T} \mu(t_i) [\lambda f - 1] < 0. \end{aligned}$$

hence $1 \notin Z(t_i)$. □

¹⁰See Aliprantis and Border (1994, Theorem 14.91). We need only check that $\rho_i(t_i)(\cdot) [\psi(a_{-i}^0, a_{-i,1}^*, \dots, a_{-i,N_2}^*, \omega)]$ is measurable in s_{-i} and $R_{-i}^{G^*,S}(s_{-i})$ is a measurable correspondence.

4.2. Measurability. The above shows that the belief mapping on the space of rules $\mu_i : \mathcal{R}_i \rightarrow \Delta(\Delta\Omega \times \mathcal{R}_{-i})$ is well-defined. Here we show that this belief mapping is measurable. We begin with a lemma identifying some measurable subsets of rationalizable rules.

Lemma 7. *For any game $G = (A_i, u_i) \in \mathcal{G}$, for any closed subset $A' \subseteq A_i$, the subset of rationalizable rules $\{r_i \in \mathcal{R}_i : r_i(G) \subseteq A'\}$ is closed in \mathcal{R}_i .*

Proof. For closed A' , the set $\mathcal{K}A' = \{K \in \mathcal{K}A_i : K \subseteq A'\}$ is closed in $\mathcal{K}A_i$ (see Aliprantis and Border (1994, Theorem 3.63)). Thus, by the definition of the product topology on \mathcal{R}_i ,

$$\{r_i \in \mathcal{R}_i : r_i(G) \subseteq A'\} = \{r_i \in \mathcal{R}_i : r_i(G) \in \mathcal{K}A'\}$$

is closed. □

Proof of Proposition 4. Let \mathcal{D} be the collection of all subsets E of $\Delta\Omega \times \mathcal{R}_{-i}$ such that the mapping

$$r_i \rightarrow \mu_i^{\mathcal{R}}(r_i)[1_E]$$

is Borel measurable. We prove the proposition by showing that \mathcal{D} includes all measurable sets.

Let $\mathcal{P}^*(\mathcal{G})$ be the collection of all finite subsets of \mathcal{G} , and define

$$\mathcal{C} = \bigcup_{\Gamma \in \mathcal{P}^*(\mathcal{G})} \left\{ V = V^0 \times \prod_{G \in \Gamma} V^G : \begin{array}{l} V^0 \subset \Delta\Omega, V^G \subset \mathcal{K}A_{-i}^G \text{ are measurable} \\ \text{and } V^G = \mathcal{K}A_{-i}^G \text{ for all } G \notin \Gamma \end{array} \right\}$$

Note that \mathcal{C} is an algebra (closed under taking complements, finite intersections and unions) and generates the product topology and hence the σ -algebra on $\Delta\Omega \times \mathcal{R}_{-i}$. We first show that $\mathcal{C} \subset \mathcal{D}$.

Consider any element $V \in \mathcal{C}$ for which $\Gamma = \{G\}$ is a singleton. If we can show that $\{r_i : \text{marg}_{\Delta\Omega \times \mathcal{K}A_{-i}^G} \mu_i^{\mathcal{R}}(r_i)(V^0 \times V^G) \in I\}$ is a measurable set of rules for every interval $I \subset [0, 1]$, it will follow that $V \in \mathcal{D}$. Since $\{\mu \in \Delta(\Delta\Omega \times \mathcal{K}A_{-i}^G) : \mu(V^0 \times V^G) \in I\}$ is a measurable set, it suffices to show that $\text{marg}_{\Delta\Omega \times \mathcal{K}A_{-i}^G} \mu_i^{\mathcal{R}} : \mathcal{R}_i \rightarrow \Delta(\Delta\Omega \times \mathcal{K}A_{-i}^G)$ is a measurable mapping.

By Lemma 6 there is a base for the Borel σ -algebra, on $\Delta\Omega \times \mathcal{K}A_{-i}^G$ consisting of sets of the form

$$W^f = \{\mu : \mu[f] < 0\}$$

for all functions $f \in \mathcal{L}^G$. In the course of the proof of Theorem 3, we showed that for any $f^G \in \mathcal{F}^G$ and $s \in \mathbf{R}$, there is a game \tilde{G} and a closed subset of actions $\tilde{A} \in A_{-i}$ such that for any type space $T \in \mathcal{TS}(\Omega)$ the following two sets of types $t_i \in T_i$ are equal

$$\{t_i : R_i^{\tilde{G}, T}(t_i) \subseteq \tilde{A}\} = \{t_i : \Delta R_i^{G, T}(t_i)[f^G] < s\}.$$

This implies that the following two sets of rationalizable rules are equal

$$\left\{ r_i \in \mathcal{R}_i : r_i(\tilde{G}) \subseteq \tilde{A} \right\} = \left\{ r_i \in \mathcal{R}_i : \operatorname{marg}_{\Delta\Omega \times \mathcal{K}A_{-i}^G} \mu_i^{\mathcal{R}}(r_i)[f^G] < s \right\}.$$

By Lemma 7, the first set is measurable. Now since $f = f_1^G - f_2^G$ for some $f_1^G, f_2^G \in \mathcal{F}^G$, we have

$$W^f = \bigcup_{\substack{x_1, x_2 \in \mathbb{Q}, \\ x_1 + x_2 < 0}} \bigcap_{m=1,2} \{ \mu : \mu[f_m^G] < x_m \},$$

where \mathbb{Q} is the set of rational numbers. Then

$$\left[\operatorname{marg}_{\Delta\Omega \times \mathcal{K}A_{-i}^G} \mu_i^{\mathcal{R}} \right]^{-1} (W^f) = \bigcup_{\substack{x_1, x_2 \in \mathbb{Q}, \\ x_1 + x_2 < 0}} \bigcap_{m=1,2} \left\{ r_i \in \mathcal{R}_i : \operatorname{marg}_{\Delta\Omega \times \mathcal{K}A_{-i}^G} \mu_i^{\mathcal{R}}(r_i)[f_m^G] < x_m \right\}$$

is measurable as countable union of finite intersections of measurable sets. Therefore the inverse image of every set in a base for the sigma-algebra is measurable and this implies that $\operatorname{marg}_{\Delta\Omega \times \mathcal{K}A_{-i}^G} \mu_i^{\mathcal{R}}$ is measurable. (See Aliprantis and Border (1994, Lemma 8.16).) Now consider an element $V \in \mathcal{C}$ for which Γ is an arbitrary finite set. Consider the product game, $\tilde{G} = \prod_{\Gamma} G$, where the product set $V^{\tilde{G}} = \prod_{G \in \Gamma} V^G$ is a measurable subset of $\mathcal{K}A_{-i}^{\tilde{G}}$. By the product structure, for any rationalizable rule $r_{-i}(\tilde{G}) = \prod_{G \in \Gamma} r_{-i}(G)$. Thus if we define $V' = V^0 \times V^{\tilde{G}} \times \prod_{G \notin \Gamma} \mathcal{K}A_{-i}^G$, we have $1_{V'} = 1_V$, and we have already shown that V' belongs to \mathcal{D} .

We have shown $\mathcal{C} \subset \mathcal{D}$. Now consider any sequence of measurable subsets $E_n \subset \Delta\Omega \times \mathcal{R}_{-i}$ such that $E_n \subset E_{n+1}$, $E \in \mathcal{D}$, and let $E = \cup E_n$. The sequence of indicator functions 1_{E_n} increases pointwise to 1_E . By countable additivity, $\mu_i^{\mathcal{R}}(r_i)(E) = \lim \mu_i^{\mathcal{R}}(r_i)(E_n)$, and hence for any open interval I ,

$$\{ r_i : \mu_i^{\mathcal{R}}(r_i)[1_E] \in I \} = \bigcup_n \bigcap_{m > n} \{ r_i : \mu_i^{\mathcal{R}}(r_i)[1_{E_m}] \in I \}$$

which is measurable. Thus $E \in \mathcal{D}$ and \mathcal{D} is a monotone class that includes the algebra \mathcal{C} . By the monotone class lemma, \mathcal{D} includes all Borel sets. \square

5. COMPACT AND FINITE GAMES

In this section we characterize rationalizable correspondences for compact games with continuous payoffs. For this class of games, the procedure of iterative elimination of never-interim-best-replies leads to $R_i^{G,T}$ in at most a countable number of steps.¹¹

¹¹Without any assumptions on the game or the type space, one can show that the iterative procedure leads to $R_i^{G,T}$ after sufficiently (transfinitely) many steps of elimination. This follows from a straightforward

The proof of Proposition 2 proceeds in two steps. First, we show that the rationalizable correspondence is non-empty and closed if the type space is a Polish space and satisfies the additional property that the mapping from types to beliefs is continuous in a strong sense. This implies in particular that the rationalizable correspondence is measurable for these type spaces and that rationalizable sets for given type are compact. The space $L(\Omega)$ introduced in section 3.6 naturally satisfies this condition.

Next, by lemma 4, from any type space, there is a measurable mapping preserving rationalizable sets to $L(\Omega)$. It will follow that the thesis of the Proposition holds for all type spaces in $\mathcal{TS}(\Omega)$.

Suppose that T is a type space and each T_i is a Polish space. Say that T is *continuous* if for each i , the mapping $\mu_i^T : T_i \rightarrow \Delta(T_{-i} \times A_{-i})$ is continuous. If in addition, there is a version of the conditional belief mapping ρ_i that is continuous then we say that T is Δ -*continuous*. Note that Δ -continuity is a stronger property than continuity alone. We have two lemmas, proved in the Appendix.

Lemma 8. *If T is Δ -continuous, and G is a compact, continuous game, then π_i is jointly continuous.*

It can be shown by an argument directly parallel to the proof of the previous lemma that U_i is jointly continuous.

If $B_{-i} \subset T_{-i} \times A_{-i}$ is an measurable assessment for i , then we say that σ_{-i}^Δ is compatible with B if $\sigma_{-i}^\Delta(B) = 1$. Given a type space, for any measurable subset $B \subset T_{-i} \times A_{-i}$, and type t_i , we let $\Sigma^\Delta(B|t_i)$ denote the set of conjectures for t_i that are concentrated on B . That is

$$\Sigma^\Delta(B|t_i) = \left\{ \sigma_{-i}^\Delta : \underset{T_{-i}}{\text{marg}} \sigma_{-i}^\Delta = \underset{T_{-i}}{\text{marg}} \mu_i(t_i) \text{ and } \sigma_{-i}^\Delta(B) = 1 \right\}$$

Lemma 9. *Suppose T is a continuous type space and let $S_i \subset T_i$ be a compact subset of types and $B \subset A \times_{-i} T_{-i}$ is a closed assessment. Then the correspondence*

$$\Sigma^\Delta(B|\cdot) : S_i \rightrightarrows \Delta(T_{-i} \times A_{-i})$$

has compact graph.

Finally, we can show the proposition:

extension of the arguments in Lipman (1994). However, in general this set may be empty and the correspondence need not be closed or even closed-valued. Compactness and continuity are used here only to deliver these properties.

Proposition 6. *Suppose that T is Δ -continuous and G is a compact game. Then the rationalizable correspondence is closed and for any type $t_i \in T_i$ in any type space $T \in \mathcal{TS}(\Omega)$,*

$$R_i^{G,T} = \mathcal{R}_i^{G,T} := \bigcap_{k \geq 0} R_{k,i}^{G,T}$$

Proof. We start with showing inductively that each $R_{n,i}^{G,T}$ is a closed correspondence. It is obviously true for $n = 0$. Suppose now it is true for some arbitrary n , and let $(t_i^k, a_i^k) \rightarrow (t_i, a_i)$ with $(t_i^k, a_i^k) \in R_{n,i}^{G,T}$. Then for each k there is a conjecture $\sigma_{-i}^{\Delta,k} \in \Sigma^\Delta(R_{n-1,-i}^{G,T}|t_i^k)$ such that $a_i^k \in B_i(\sigma_{-i}^{\Delta,k})$. By Lemma 9, there is a subsequence $\sigma_{-i}^{\Delta,l}$ converging to $\sigma_{-i}^\Delta \in \Sigma^\Delta(R_{n-1,-i}^{G,T}|t_i)$.

Finally, by the continuity of $U_i(a_i, \sigma_{-i}^\Delta)$,

$$U_i(a_i^l, \sigma_{-i}^{\Delta,l}) \rightarrow U_i(a_i, \sigma_{-i}^\Delta)$$

$$U_i(z_i, \sigma_{-i}^{\Delta,l}) \rightarrow U_i(z_i, \sigma_{-i}^\Delta)$$

Thus, $a_i^l \in B_i(\sigma_{-i}^{\Delta,l})$ for all l implies $a_i \in B_i(\sigma_{-i}^\Delta)$. We have shown that $(t_i, a_i) \in R_{n,i}^{G,T}$ and hence that the latter is closed.

The first step implies that correspondence $\mathcal{R}_i^{G,T}$ is closed as intersection of closed sets. Now, we move to show that $R_i^{G,T} = \mathcal{R}_i^{G,T}$. Because $R_i^{G,T}$ has the fixed-point property, we have $R_i^{G,T} \subset R_{k,i}^{G,T}$ for every k , hence $R_i^{G,T}$ is contained in $\mathcal{R}_i^{G,T}$. To show equality, therefore, it suffices to show that $\mathcal{R}_i^{G,T}$ also has the fixed-point property and is therefore a subset of $R_i^{G,T}$. We need to show

$$\mathcal{R}_i^{G,T} = \{(t_i, a_i) : a_i \in B_i(\sigma_{-i}^\Delta) \text{ for some } \sigma_{-i}^\Delta \in \Sigma^\Delta(\mathcal{R}_{-i}^{G,T}|t_i)\}.$$

Suppose $a_i \in B_i(\sigma_{-i}^\Delta)$ for some $\sigma_{-i}^\Delta \in \Sigma^\Delta(\mathcal{R}_{-i}^{G,T}|t_i)$. Then $\sigma_{-i}^\Delta \in \Sigma^\Delta(R_{k,-i}^{G,T}|t_i)$ for every k and hence $(t_i, a_i) \in R_{k,i}^{G,T}$ for every k . This shows that $(t_i, a_i) \in \mathcal{R}_i^{G,T}$.

Suppose $(t_i, a_i) \in \mathcal{R}_i^{G,T}$, i.e. $(t_i, a_i) \in R_{k,i}^{G,T}$ for every k . Then for each k there is a $\sigma_{-i}^{\Delta,k} \in \Sigma^\Delta(R_{k,-i}^{G,T}|t_i)$ such that $a_i \in B_i(\sigma_{-i}^{\Delta,k})$. Since $\mathcal{R}_i^{G,T}$ is closed, we are allowed to use lemma 9 to extract a convergent subsequence $\sigma_{-i}^{\Delta,l} \rightarrow \sigma_{-i}^\Delta$. Argument above should convince us that $a_i \in B_i(\sigma_{-i}^\Delta)$. In order to conclude that (t_i, a_i) is best response to some conjecture from $\Sigma^\Delta(\mathcal{R}_{-i}^{G,T}|t_i)$, it is enough to check that $\sigma_{-i}^\Delta \in \Sigma^\Delta(\mathcal{R}_{-i}^{G,T}|t_i)$. Notice however that this is immediate consequence of two facts:

$$\text{marg}_{T_i} \sigma_{-i}^\Delta = \lim_{l \rightarrow \infty} \text{marg}_{T_i} \sigma_{-i}^{\Delta,l} = \text{marg}_{T_i} \mu_i(t_i),$$

$$\sigma_{-i}^\Delta(\mathcal{R}_{-i}^{G,T}) = \lim_{n \rightarrow \infty} \sigma_{-i}^{\Delta,l}(R_{n,i}^{G,T}) = \lim_{n \rightarrow \infty} \lim_{l \rightarrow \infty} \sigma_{-i}^{\Delta,l}(R_{n,i}^{G,T}) = 1.$$

The last equality follows from the definition of $\sigma_{-i}^{\Delta,l} \in \Sigma^\Delta(R_{k(l),-i}^{G,T}|t_i)$ for some $k(l) : \text{for } k \leq k(l)$ we have $\sigma_{-i}^{\Delta,l}(R_{k,i}^{G,T}) = \sigma_{-i}^{\Delta,l}(R_{k(l),i}^{G,T}) = 1$. This ends the proof. \square

Proof of proposition 2. By the previous proposition, if T is Δ -continuous, then the rationalizable correspondence is closed. Because the range space A_i is compact, $R_i^{G,T}$ is upper hemi-continuous, i.e. for every closed $F \subset A_i$, the set $\{t_i \in T_i : R_i^{G,T}(t_i) \cap F \neq \emptyset\}$ is closed, and in particular measurable. Now by Corollary 14.70 in Aliprantis and Border (1994), it is measurable viewed as a function from T_i to $\mathcal{K}A_i$.

Notice that $L(\Omega)$ is Δ -continuous: conditional beliefs given type of the opponent (τ_{-i}, u_{-i}) are equal to τ_{-i} regardless of what is type of player i, t_i . Hence, rationalizable correspondence on $L(\Omega)$ is closed.

Lemma 4 shows that any type space $T \in \mathcal{TS}(\Omega)$ can be mapped by a type mapping which preserves rationalizable sets into space $L(\Omega)$ ¹². Thus, for any given type space $T \in \mathcal{TS}(\Omega)$, the rationalizable correspondence R_i^G, T is the composition of this measurable mapping and the rationalizable and closed correspondence for $L(\Omega), R_i^G, L(\Omega)$. Thus, R_i^G, T is non-empty, closed-valued, and when viewed as a function, measurable. \square

6. COMMENTS AND OTHER EXAMPLES

6.1. Example. The conventional universal type space $U(\Omega)$ is not rich enough from the point of view of solution concepts such as Bayesian equilibrium or Rationalizability. We have previously shown this by demonstrating that there are types whose rationalizable rules are not represented by any type in $U(\Omega)$. Here we present an example which makes the point even stronger: there is an action which is not rationalizable for any type in $U(\Omega)$, yet as we show below, it is easy to construct simple, perfectly standard type spaces in which the action is rationalizable. Consider the two-player game with two states of the world with payoff matrix given in figure 1.

	a_2	b_2	b'_2	a'_2
a_1	1,1	1,-9	-1,-9	-1,-1
b_1	-9,1	0,0	-9,-9	-9,-1
b'_1	-9,-1	-9,-9	0,0	-9,1
a'_1	-1,-1	-1,-9	1,-9	1,1

$\omega = +1$

	a_2	b_2	b'_2	a'_2
a_1	1,1	1,-9	-1,-9	-1,-1
b_1	-9,1	-9,-9	0,0	-9,-1
b'_1	-9,-1	0,0	-9,-9	-9,1
a'_1	-1,-1	-1,-9	1,-9	1,1

$\omega = -1$

FIGURE 1. Actions b_i and b'_i are not rationalizable in $U(\Omega)$.

¹²Note that through the proof of this lemma we do not need results from section 2.4. In other words, there is no circularity in the argument.

We will show that neither b_i nor b'_i are rationalizable for any type in $U(\Omega)$. Note first that an equal mixture between a_i and a'_i guarantees a payoff of 0. Thus, b_i and b'_i are best-replies only if player i is certain that the opponent plays an action in $\{b_{-i}, b'_{-i}\}$, and the action is correlated with the state. Now if i assigns greater than 1/2 probability to state +1, then it is easily verified that action a_i achieves strictly higher payoff than b_i , and action a'_i achieves strictly higher payoff than b'_i , regardless of the opponent's strategy. Likewise, if the probability of state +1 is less than 1/2, then a_i must do better than b'_i and a'_i better than b_i . Thus, actions b_i or b'_i can be rationalizable only for types who assign the two states equal probability and who assign probability 1 to opponent's types for whom b_{-i} or b'_{-i} are rationalizable. Now the game is symmetric, so the same analysis applies to player $-i$ with the player's roles reversed. Putting these two conclusions together, actions b_i and b'_i are rationalizable only for types of player i who assign equal probability to the two states, and probability 1 to the event that player $-i$ has the same beliefs and assigns probability 1 to the event that b_i and b'_i are rationalizable for i . By induction, b_i and b_{-i} are rationalizable only for those types of player i for whom it is common-knowledge that the two states are equally likely. Let v_i be the type in $U_i(\Omega)$ with this hierarchy of beliefs, and v_{-i} the analogous type for player $-i$. Note that in $U(\Omega)$, type v_i assigns probability 1 to v_{-i} and equal probability to the two states. But then no matter what mixed action is played by v_{-i} , it is never correlated with the state. Thus b_i and b'_i can never be best-replies, hence never be rationalizable for type v_i .

Nevertheless, both b_i and b'_i are rationalizable for all types in the type space from the introduction. Indeed, any pure strategy profile in which the two types of each player play different actions in $\{b_i, b'_i\}$ is a Bayesian Nash equilibrium.

6.2. Upper-Hemicontinuity. The literature has had some interest in finding the “right” topology on the universal type space to capture similarity of types with respect to their strategic behavior. One requirement of such a topology is that the rationalizable correspondence should be upper hemi-continuous. Our results shed some light on what would be required of such a topology. We have shown (Proposition 6) that a sufficient condition for upper hemi-continuity is that the topology be fine enough so that conditional beliefs are continuous. Here we show that this is in general necessary as well.

Suppose that $\Omega = \{-1, 1\}$ and consider the game from the introduction and the following type space: $T_i = [-1, 1]$ and beliefs are defined by

$$\mu_i(t_i) \{t_{-i}, \omega\} = \begin{cases} \frac{1}{2} & \text{if } t_{-i} \in \{t_i, -t_i\} \text{ and } \text{sign}[t_i \cdot \omega] = \text{sign}[t_{-i}], \\ 0 & \text{otherwise} \end{cases}$$

It is easy to verify that these beliefs are generated by a common prior and they are continuous as a function of t_i . However conditional beliefs over Ω exhibit a discontinuity at $t_i = t_{-i} = 0$. Indeed, if $t_i = t_{-i} \neq 0$, then t_i assigns probability 1 to state $\omega = 1$ conditional on t_{-i} , but if $t_i = t_{-i} = 0$, then both states have equal conditional probability.

For $t_i = t_{-i} \neq 0$, the set $\{-t_i, +t_i\} \times \{-t_{-i}, t_{-i}\}$ is a belief-closed subspace which is isomorphic to the first type space from the introduction. Thus, all actions are rationalizable for these types. However, the belief-closed subspace $\{0\} \times \{0\}$, is isomorphic to the second type space and hence action c_i is the unique rationalizable action for types $t_i = t_{-i} = 0$. Thus, for this finite game with a finite set of states of the world and continuous belief-mapping, the rationalizable correspondence is not upper hemi-continuous.

6.3. Discontinuous Games. We have defined rationalizability as a fixed-point of the interim best-reply operator. Just as Bernheim (1984) and Pearce (1984) showed in the complete information case, this is equivalent to the more customary iterative definition for compact and continuous games. However, the fixed point definition is more flexible as it is guaranteed to exist (although possibly empty) for any game (note that the proof of proposition 1 assumes nothing about the game.) Thus, all of our results that do not rely on the iterative characterization of rationalizability in extend immediately to discontinuous games. For example, Lemma 3 uses only the fixed-point definition and since the conclusion of Proposition 5 relies only on Lemma 3, we can extend our results to conclude that for *any* game, two types have the same rationalizable sets if they have the same Δ -hierarchies. In particular this is true even for games in which the rationalizable sets are reached only after transfinitely many rounds of elimination, see Lipman (1994).

6.4. Universal Type Space for the Measurable Case. Following the literature, we say that a type space U over a space of basic uncertainty X is *universal* among type spaces with property Y if for every such type space there is a unique mapping into U which preserves beliefs. Mertens, Sorin, and Zamir (1994) showed that there exists a universal type space for all continuous type spaces, assuming X is a Polish space. On the other hand, Heifetz and Samet (1999) showed that there is no universal type space for measurable (not necessarily continuous) type spaces when X is assumed only to be measurable. Our Theorem 1 is a positive result for an in-between case. It shows the existence of a universal measurable type space under the assumption that X is Polish. This may be comforting because while there may be good reason to assume some structure on the physical world X , but there is no good reason to assume structure on a type space which is nothing more than an artificial modeling construct.

REFERENCES

- ALIPRANTIS, C. D., AND K. C. BORDER (1994): *Infinite Dimension Analysis: A Hitchhiker's Guide*. Springer-Verlag, New York.
- AUMANN, R. J. (1987): "Correlated Equilibrium as an Expression of Bayesian Rationality," *Econometrica*, 55, 1–18.
- BATTIGALLI, P., AND M. SINISCALCHI (1999): "Hierarchies of Conditional Beliefs and Interactive Epistemology in Dynamic Games," *Journal of Economic Theory*, 88, 188–230.
- BERNHEIM, D. (1984): "Rationalizable Strategic Behavior," *Econometrica*, 52, 1007–1028.
- BRANDENBURGER, A., AND E. DEKEL (1987): "Rationalizability and Correlated Equilibria," *Econometrica*, 55, 1391–1402.
- (1993): "Hierarchies of Beliefs and Common Knowledge," *Journal of Economic Theory*, 59, 189–198.
- DEKEL, E., D. FUDENBERG, AND S. MORRIS (2003): "Interim Rationalizability," mimeo.
- GUL, F. (1998): "A Comment on Aumann's Bayesian View," *Econometrica*, 66, 923–927.
- HARSANYI, J. C. (1967-68): "Games with Incomplete Information Played by Bayesian Players, I, II, III," *Management Science*, 14, 159–182, 320–334, 486–502.
- HEIFETZ, A., AND D. SAMET (1999): "Coherent Beliefs are not Always Types," *Journal of Mathematical Economics*, 32, 475–488.
- LIPMAN, B. (1994): "A Note on the Implications of Common Knowledge of Rationality," *Games and Economic Behavior*, 6, 114–129.
- MERTENS, J.-F., S. SORIN, AND S. ZAMIR (1994): "Repeated Games," CORE Discussion Papers 9420,9421,9422; Louvain-la-Neuve.
- MERTENS, J. F., AND S. ZAMIR (1985): "Formulation of Bayesian Analysis for Games with Incomplete Information," *International Journal of Game Theory*, 14, 1–29.
- MORRIS, S. (2002): "Typical Types," <http://www.econ.yale.edu/sm326/typical.pdf>.
- PEARCE, D. (1984): "Rationalizable Strategic Behavior and the Problem of Perfection," *Econometrica*, 52, 1029–1050.

APPENDIX A. SKETCH OF PROOF OF THEOREM 1

We follow a construction of universal type space over Polish space X from Mertens, Sorin, and Zamir (1994). Suppose that we have a sequence of Polish spaces $\{P_n\}_{n \geq 0}$ with a sequence of continuous mappings $h_n : P_{n+1} \rightarrow P_n$. Then the projective limit is a subset $P \subseteq P_0 \times P_1 \times \dots$ of all points (p_0, p_1, \dots) , such that $h_n(p_{n+1}) = p_n$ for all n . We can show that such set is a Polish space in product topology. There are continuous induced mappings $H_n : P \rightarrow P_n$, which are projections of sequence on its n th coordinate.

We may construct an universal space $U(X)$ for Polish X as a projective limit of hierarchies of beliefs. Define inductively spaces of n th hierarchy of beliefs of player i - U_i^n as

$$U_i^0 = \{*\} \text{ and for } n > 0 \ U_i^n = \Delta(X \times U_{-i}^{n-1}).$$

Define inductively mappings $h_i^n : U_i^{n+1} \rightarrow U_i^n$ with

$$h_i^0(u_i) = * \text{ for } u_i \in U_1 \text{ and}$$

$$h_n = \Delta \left(\text{id}_X \times h_{-i}^{n-1} \right) \text{ for } n > 0.$$

The last mapping is a transported measure through mapping $\text{id}_X \times h_{-i}^{n-1} : X \times U_{-i}^{n-1} \rightarrow X \times U_{-i}^{n-2}$. Then $U(X)$ is an projective limit of the system $(U_i^n, h_i^n)_i$. There are induced mappings $H_i^n : U_i^X \rightarrow U_i^n$, such that

$$H_i^n = h_i^n \circ H_i^{n+1}.$$

Mertens, Sorin, and Zamir (1994) show that if type space T is continuous (belief mapping is continuous) then there exists a sequence of continuous $u_i^{T,n} : T_i \rightarrow U_i^n$, $h_i^n \circ u_i^{T,n+1} = u_i^{T,n}$, which extend to the continuous mapping $u_i^T : T_i \rightarrow U_i$, which is a unique exact and preserves beliefs. The difference in our case is that we are not able to guarantee continuity of mappings $u_i^{T,n}$. However, we show that weak measurability of belief mapping assures that maps $u_i^{T,n}$ defined exactly in the same way as in Mertens, Sorin, and Zamir (1994) are measurable. Moreover, the generate pointwise converging measurable mappings $u_i^{T,n,u_i^*} : T_i \rightarrow U_i$, for some $u_i^* \in U_i$, which converge to to measurable mapping $u_i^T : T_i \rightarrow U_i$. This mapping is, by the same reasons, unique exact and belief preserving mapping.

Precisely, we use the following lemma:

Lemma 10. *Suppose that there is a Polish space B and measurable mapping $\phi_{-i} : T_{-i} \rightarrow B$. Then mapping $\Phi_i : T_i \rightarrow \Delta(X \times B)$ defined with*

$$\Phi_i(t_i) = \Delta \left(\text{id}_X \times \phi_{-i} \right) \mu_i(t_i)$$

is measurable.

Proof. We need to check whether for any measurable function $f : X \times B \rightarrow R$, sets $\{t_i : \Phi_i(t_i)[f] < 0\}$ are measurable. But

$$\{t_i : \Phi_i(t_i)[f] < 0\} = \{t_i : \mu_i(t_i)[f(x, \phi_{-i}(t_{-i}))] < 0\}$$

and the last set is measurable due to definition of weak measurability of belief mapping. \square

Choose now some $u_i^* \in U_i$ and construct mappings $u_i^{T,n,u_i^*} : T_i \rightarrow U_i$ with $u_i^{T,n,u_i^*}(t_i) = u_i^*$ and later inductively

$$u_i^{n+1} = \Delta \left(\text{id}_X \times u_{-i}^n \right) \mu_i(t_i).$$

By lemma, each of these mappings is measurable. Moreover, they converge pointwise to mapping u_i^T , which is also measurable (as a pointwise limit of measurable mappings). We check as in Mertens, Sorin, and Zamir (1994) that it is unique exact and belief-preserving mapping.

APPENDIX B. PROOF OF LEMMA 2

We need the following result

Lemma 11. *Let A and B be measurable spaces and $g : A \times B \rightarrow [0, 1]$ a jointly measurable map. If $m : A \rightarrow \Delta(B)$ is measurable, then the map $L^g : A \rightarrow \mathbf{R}$ defined by $L^g(a) = m(a)[g(a, \cdot)]$ is measurable.*

Proof. There exists a sequence of simple functions $g_n : A \times B \rightarrow \mathbf{R}$ such that $g_n \uparrow g$ and by the definition of the Lebesgue integral, for any probability measure $\nu \in \Delta(A \times B)$.

$$\nu[g_n] \rightarrow \nu[g]$$

In particular, for any given $a \in A$, if ν is the measure whose marginal on B is $m(a)$ and whose marginal on A is δ_a ,

$$L^{g_n}(a) = \nu[g_n] \rightarrow \nu[g] = L^g(a)$$

Thus, if we can show that L^f is measurable for all simple functions f , then we will have shown that L^g is measurable as the pointwise limit of measurable mappings.

First consider $f = 1_{\alpha \times \beta}$ for $\alpha \subset A$ and $\beta \subset B$ (measurable.) We have $L^f(a) = 1_\alpha(a) \cdot m(a)(\beta)$ which is measurable since m was assumed to be measurable. Thus, L^f is measurable for all f that are indicators of product sets. Now for any finite k , let $\alpha_1, \dots, \alpha_k$ and β_1, \dots, β_k be measurable subsets of A and B respectively and note that for $f = \prod_{l=1}^k 1_{\alpha_l \times \beta_l}$,

$$L^f(a) = \prod_l 1_{\alpha_l} \cdot m(a)(\cap_l \beta_l)$$

is measurable. Thus if $f = 1_{\cap_l (\alpha_l \times \beta_l)} = \prod_l 1_{\alpha_l \times \beta_l}$, then L^f is measurable, and if

$$f = 1_{\cup_l (\alpha_l \times \beta_l)} = \sum_l 1_{\alpha_l \times \beta_l} - \sum_{S \subset \{1, \dots, k\}} (|S| - 1) \prod_{l \in S} 1_{\alpha_l \times \beta_l}$$

then L^f is measurable as a linear combination of measurable functions. Note also that $L^{1-E} = L^{1-E} = 1 - L^{1E}$. Thus L^f is measurable for all indicator functions f of sets in the algebra generated by the product sets.

Now consider any sequence X_n with $X_n \subset A \times B$, $X_n \subset X_{n+1}$ for all n and $\cup_n X_n = X$. The corresponding sequence L^{1X_n} is an increasing sequence of maps converging pointwise to L^{1X} . Thus if L^{1X_n} are measurable for all n , so is L^{1X} . It follows that the collection of sets X such that L^{1X} is measurable is a monotone class. Since it includes the algebra generated by the product sets, by the monotone class lemma it includes the corresponding σ -algebra,

i.e. the product sigma-algebra on $A \times B$. Finally, since any simple function $f : A \times B \rightarrow \mathbf{R}$ has the form

$$f(a, b) = \sum_{j=1}^k c_j 1_{X_j}(a, b)$$

for some coefficients c_j and measurable sets $X_j \subset A \times B$, any such L^f is measurable as a linear combination of measurable functions. \square

Proof of Lemma 2. We must show that for any measurable $f : \Delta\Omega \times T_{-i} \rightarrow \mathbf{R}$, the mapping

$$t_i \rightarrow \mu_i^\Delta(t_i)[f]$$

is measurable. Define $g(t_i, t_{-i}) = f(\rho_i(t_i, t_{-i}), t_{-i})$. Note that g is jointly measurable and

$$\mu_i^\Delta(t_i)[f] = \mu_i(t_i)[g(t_i, \cdot)].$$

Now apply lemma 11. \square

APPENDIX C. PROOF OF LEMMA 6

Let \mathbf{H} denote the Hilbert cube $[0, 1]^\mathbb{N}$. Since Ω is Polish, there is a countable sequence of functions $h_k^* : \Omega \rightarrow [0, 1]$, which define a compatible metric on $\Delta\Omega$, $d_{\Delta\Omega}(\tau, \tau') = \sum_{k=1}^{\infty} \frac{1}{2^k} |\tau[h_k^*] - \tau'[h_k^*]|$ and a mapping $H : \Delta\Omega \rightarrow \mathbf{H}$

$$H(\tau) = (\tau[h_1^*], \tau[h_2^*], \dots).$$

The mapping H embeds $\Delta\Omega$ (with the weak \blacksquare topology) in the Hilbert cube (with the product topology). Suppose we have a family \mathcal{F} of continuous functions $f : \mathbf{H} \times \mathcal{K}A_{-i} \rightarrow \mathbf{R}$ such that the collection of sets

$$\{\mu : \mu[f(h, A)] < 0\} \subseteq \Delta(\mathbf{H} \times \mathcal{K}A_{-i})$$

is a subbase for the weak \blacksquare topology on $\Delta(\mathbf{H} \times \mathcal{K}A_{-i})$ (in which case we say that \mathcal{F} generates the topology). Then, because H is an embedding, it will follow that the corresponding family $\mathcal{F}' \subset C(\Delta\Omega \times \mathcal{K}A_{-i})$

$$\mathcal{F}' = \{f' : f'(\tau, K) = f(H(\tau), K) \text{ for some } f \in \mathcal{F}\}$$

generates the topology on $\Delta(\Delta\Omega \times \mathcal{K}A_{-i})$. The strategy of proof is to find such an \mathcal{F} so that the corresponding \mathcal{F}' is included in \mathcal{L}^G .

For each natural number n , define the following set of continuous functions $f : [0, 1]^n \times \mathcal{K}A_{-i} \rightarrow [0, \infty)$:

$$\mathcal{F}_n = \left\{ \begin{array}{l} f(z_1, \dots, z_n, A) = \max_{k=1, \dots, N_1} \sup_{a_1, \dots, a_{N_2} \in A} \eta^k(a_1, \dots, a_{N_2}) \cdot z : \\ \text{for some natural numbers } N_1, N_2 \text{ and} \\ \text{some continuous bounded functions } \eta^1, \dots, \eta^{N_1} : A^{N_2} \rightarrow [0, 1]^n. \end{array} \right\}$$

where \cdot is a scalar product of two vectors in \mathbf{R}^n . Next, define set of differences of functions from \mathcal{F}_n

$$\mathcal{L}_n = \{f - g : f, g \in \mathcal{F}_n\} \subseteq C([0, 1]^n \times \mathcal{K}A_{-i}).$$

We have a lemma:

Lemma 12. *Set \mathcal{L}_n is uniformly dense in the set $C([0, 1]^n \times \mathcal{K}A_{-i})$.*

Proof. This is a standard argument applying the lattice version of the Stone-Weierstrass theorem (see Aliprantis and Border (1994, Theorem 7.45)). We need to verify that \mathcal{L}_n :

- is closed under scalar multiplication: If $(f-g) \in \mathcal{L}_n$, then for any $\lambda \in \mathbf{R}$, $\lambda(f-g) \in \mathcal{L}_n$ as well;
- contains constant function: $\mathbf{1} \in \mathcal{L}_n$;
- is closed under finite sums: first note that for any $f, g \in \mathcal{F}_n$, $z = (z_1, \dots, z_m) \in [0, 1]^n$ and $K \in \mathcal{K}A_{-i}$

$$\begin{aligned} & f(z, K) + g(z, K) \\ &= \max_{k=1, \dots, N_1^f} \sup_{a_1, \dots, a_{N_2^f} \in K} \eta^k(a_1, \dots, a_{N_2^f}) \cdot z \\ &+ \max_{l=1, \dots, N_1^g} \sup_{a_1, \dots, a_{N_2^g} \in K} \nu^l(a_1, \dots, a_{N_2^g}) \cdot z \\ &= \max_{\substack{k=1, \dots, N_1^f \\ l=1, \dots, N_1^g}} \sup_{\substack{a_1, \dots, a_{N_2^f} \in K \\ a_1, \dots, a_{N_2^g} \in K}} \left(\eta^k(a_1, \dots, a_{N_2^f}) + \nu^l(a_1, \dots, a_{N_2^g}) \right) \cdot z \end{aligned}$$

so that $f + g \in \mathcal{F}_n$. But this implies that for any $(f - g), (f' - g') \in \mathcal{L}_n$ we also have $(f + f') - (g + g') \in \mathcal{L}_n$;

- is closed with respect to taking maximum of two functions: for any $f, g \in \mathcal{F}_n$, $z \in [0, 1]^n$, $K \in \mathcal{KA}_{-i}$

$$\begin{aligned}
& \max \{f(z, A), g(z, A)\} \\
&= \max \left\{ \begin{array}{l} \max_{k=1, \dots, N_1^f} \sup_{a_1, \dots, a_{N_2^f} \in K} \eta^k(a_1, \dots, a_{N_2^f}) \cdot z, \\ \max_{l=1, \dots, N_1^g} \sup_{a_1, \dots, a_{N_2^g} \in K} \nu^l(a_1, \dots, a_{N_2^g}) \cdot z \end{array} \right\} \\
&= \max_{k=1, \dots, N_1^f + N_1^g} \sup_{a_1, \dots, a_{N_2^f + N_2^g} \in K} \varphi^k(a_1, \dots, a_{N_2^f + N_2^g}) \cdot z
\end{aligned}$$

where

$$\begin{aligned}
\varphi^k(a_1, \dots, a_{N_2^f + N_2^g}) &= \eta^k(a_1, \dots, a_{N_2^f}) \text{ for } k \leq N_1^f \text{ and} \\
\varphi^k(a_1, \dots, a_{N_2^f + N_2^g}) &= \nu^{k - N_1^f}(a_{N_2^f + 1}, \dots, a_{N_2^f + N_2^g}) \text{ for } N_1^f < k \leq N_1^f + N_1^g.
\end{aligned}$$

Then $h = \max(f, g) \in \mathcal{F}_n$. Together with the fact that

$$\max \{f - g, f' - g'\} = \max \{f + g', f' + g\} - (g + g')$$

and the previous point, it implies that $\max \{f - g, f' - g'\} \in \mathcal{L}_n$ for any $f - g, f' - g' \in \mathcal{L}_n$;

- separates points: for any $z, z' \in [0, 1]^n$, $z \neq z'$, there is vector $\eta \in \mathbf{R}^n$, such that $\eta \cdot z \neq \eta \cdot z'$. Similarly, for any $K, K' \in \mathcal{KA}_{-i}$, $K \neq K'$, there is a continuous function $s : A_{-i} \rightarrow [0, 1]$, such that

$$f(A) = \sup_{a \in K} s(a) = 1 > 0 = \sup_{a \in K'} s(a) = f(A').$$

□

Finally we can prove lemma 6. Any $f \in C([0, 1]^n \times \mathcal{KA}_{-i})$ can be viewed as an element $f' \in C(\mathbf{H} \times \mathcal{KA}_{-i})$ by writing $f'(h, K) = f(h_1, \dots, h_n, K)$. By the Stone-Weierstrass theorem (algebraic version, see Aliprantis and Border (1994, Theorem 7.46)) $\cup_n C([0, 1]^n \times \mathcal{KA}_{-i})$ is uniformly dense in $C(\mathbf{H} \times \mathcal{KA}_{-i})$. By lemma 12, family \mathcal{L}_n is uniformly dense in $C([0, 1]^n \times \mathcal{KA}_{-i})$. Thus $\cup_n \mathcal{L}_n$ is uniformly dense in $\cup_n C([0, 1]^n \times \mathcal{KA}_{-i})$ and hence in $C(\mathbf{H} \times \mathcal{KA}_{-i})$. We conclude that the family $\cup_n \mathcal{L}_n$ generates the topology on $\Delta(\mathbf{H} \times \mathcal{KA}_{-i})$ (see Aliprantis and Border (1994, Theorem 12.2)).

The proof is now completed by showing that each $f \in \mathcal{L}_n$ corresponds to a function f' belonging to \mathcal{L}^G by the formula $f'(\tau, K) = f(H(\tau), K)$. By the linear structure of \mathcal{L}^G it suffices to show that for each $g \in \mathcal{F}_n$, the composition $g \circ H : \Delta\Omega \times \mathcal{KA}_{-i}^G \rightarrow \mathbf{R}$ belongs to

\mathcal{F}^G . This is verified by noting that

$$\begin{aligned} (g \circ H)(\tau, K) &= \max_{k=1, \dots, N_1} \sup_{a_1, \dots, a_{N_2} \in K} \eta^k(a_1, \dots, a_{N_2}) \cdot (\tau[h_1^*(\omega)], \dots, \tau[h_n^*(\omega)]) \\ &= \max_{k=1, \dots, N_1} \sup_{a_1, \dots, a_{N_2} \in A} \tau \left[\sum_{m=1}^n h_m^*(\omega) \eta_m^k(a_1, \dots, a_{N_2}) \right]. \end{aligned}$$

Since $\sum_{m=1}^n h_m^*(\omega) \eta_m^k(a_1, \dots, a_{N_2})$ is a bounded continuous function from $\{1, \dots, N_1\} \times (A_{-i}^G)^{N_2} \times \Omega \rightarrow \mathbf{R}$, $g \circ H$ is an element of \mathcal{F}^G . \square

APPENDIX D. PROOFS OF LEMMAS FROM SECTION 5

Proof of Lemma 8. Pick $M > \sup |u_i(a, \omega)|$ (recall that we assume that u_i is bounded for this class of games.) Let $(a^k, t^k) \rightarrow (a^\infty, t^\infty) \in A \times T$. The set $\{t^k\}_{k=1}^\infty$ is compact, and so by Δ -continuity, the corresponding family of measures $\{\rho_i(t^k)\} \subset \Delta\Omega$ is also compact. Because Ω is a Polish space, the family is tight, i.e. for every $\varepsilon > 0$, there exists a compact $K^\varepsilon \subset \Omega$ such that $\rho_i(t^k)(K^\varepsilon) > 1 - \varepsilon$ for all $k \in \{1, \dots, \infty\}$. We have

$$|\pi_i(a^k, t^k) - \pi_i(a^\infty, t^\infty)| \leq \left| \int_{K^\varepsilon} u_i(a^k, \cdot) d\rho_i(t^k) - \int_{K^\varepsilon} u_i(a^\infty, \cdot) d\rho_i(t^\infty) \right| + 2\varepsilon M$$

Since K^ε is compact and u_i is continuous,

$$\sup_{\omega \in K^\varepsilon} |u_i(a^k, \omega) - u_i(a^\infty, \omega)| \rightarrow 0$$

i.e., the sequence of functions $u_i(a^k, \cdot) : K^\varepsilon \rightarrow \mathbf{R}$ converges uniformly to $u_i(a^\infty, \cdot)$. It follows that

$$\left| \int_{K^\varepsilon} u_i(a^k, \omega) d\rho_i(t^k) - \int_{K^\varepsilon} u_i(a^\infty, \cdot) d\rho_i(t^\infty) \right| \rightarrow 0$$

and so

$$\begin{aligned} \limsup_k |\pi_i(a^k, t^k) - \pi_i(a^\infty, t^\infty)| &\leq \limsup_k \left| \int_{K^\varepsilon} u_i(a^k, \cdot) d\rho_i(t^k) - \int_{K^\varepsilon} u_i(a^\infty, \cdot) d\rho_i(t^\infty) \right| + 2\varepsilon M \\ &= 2\varepsilon M \end{aligned}$$

and since ε was arbitrary, we have shown $\pi_i(a^k, t^k) \rightarrow \pi_i(a^\infty, t^\infty)$. \square

Proof of Lemma 9. The proof uses the following result (see Aliprantis and Border (1994, Theorem 12.20)): If X is a Polish space, then a family $\mathcal{F} \subset \Delta(X)$ has compact closure if and only if \mathcal{F} is tight, i.e. for every $\varepsilon > 0$ there is a compact $K \subset X$ such that $\nu(K) > 1 - \varepsilon$ for all $\nu \in \mathcal{F}$.

Pick $\varepsilon > 0$. Since S_i is compact, by the continuity of T so is $\mu_i^T(S_i) = \{\mu_i^T(t_i) : t_i \in S_i\}$ and by the continuity of marginals, so is $\text{marg}_{T_{-i}} \mu_i^T(S_i)$. By the above result, there is a compact $K \subset T_{-i}$ such that $\text{marg}_{T_{-i}} \mu_i^T(t_i)(K) > 1 - \varepsilon$ for all $t_i \in S_i$. Thus for any $t_i \in S_i$

and $\sigma_{-i}^\Delta \in \Sigma^\Delta(B|t_i)$, we have $\sigma_{-i}^{\Delta, t_i}(K \times A_{-i}) = \text{marg}_{T_{-i}} \mu_i^T(t_i)(K) > 1 - \varepsilon$. Since $K \times A_{-i}$ is compact, this shows that the family

$$\bigcup_{t_i \in S_i} \Sigma^\Delta(B|t_i)$$

is tight and therefore has compact closure.

Because $\sigma_{-i}^{\Delta, k}$ is a sequence from this set, it has a convergent subsequence $\sigma_{-i}^{\Delta, l} \rightarrow \sigma_{-i}^\Delta$. The proof is concluded by showing that $\sigma_{-i}^\Delta \in \Sigma^\Delta(B|t_i)$.

- (1) (δ_x denotes Dirac measure on x .) $\text{marg}_{T_i} \sigma_{-i}^{\Delta, l} = \delta_{t_i^l} \rightarrow \delta_{t_i}$ and by continuity of marginals, $\text{marg}_{T_i} \sigma_{-i}^{\Delta, l} \rightarrow \text{marg}_{T_i} \sigma_{-i}^\Delta$, thus $\text{marg}_{T_i} \sigma_{-i}^\Delta = \delta_{t_i}$.
- (2) $\text{marg}_{T_{-i}} \sigma_{-i}^{\Delta, l} = \text{marg}_{T_{-i}} \mu_i^T(t_i^l)$. Because T is continuous, $\text{marg}_{T_{-i}} \mu_i^T(t_i^l) \rightarrow \text{marg}_{T_{-i}} \mu_i^T(t_i)$ and since $\text{marg}_{T_{-i}} \sigma_{-i}^{\Delta, l} \rightarrow \text{marg}_{T_{-i}} \sigma_{-i}^\Delta$, we have $\text{marg}_{T_{-i}} \sigma_{-i}^\Delta = \text{marg}_{T_{-i}} \mu_i^T(t_i)$.
- (3) $M_{T_{-i} \times A_{-i}} \sigma_{-i}^{\Delta, l}(B) = 1$. Since taking marginals is continuous and the set of probability measures assigning probability 1 to a closed set is closed, $M_{T_{-i} \times A_{-i}} \sigma_{-i}^\Delta(B) = M_{T_{-i} \times A_{-i}}[\lim \sigma_{-i}^{\Delta, l}](B) = 1$.

□