Uniform Topology on Types and Strategic Convergence*

Alfredo Di Tillio IGIER and IEP Università Bocconi alfredo.ditillio@unibocconi.it Eduardo Faingold Department of Economics Yale University eduardo.faingold@yale.edu

August 17, 2007 (Preliminary draft. Comments welcome.)

Abstract

We study the continuity of the correspondence of interim ε -rationalizable actions in incomplete information games. We introduce a topology on types, called *uniform-weak topology*, under which two types of a player are close if they have similar first-order beliefs, attach similar probabilities to other players having similar first-order beliefs, and so on, where the degree of similarity is uniform over the levels of the belief hierarchy. This notion of proximity of types is an extension of the concept of common *p*-belief due to Monderer and Samet (1989). We show that, given any finite game, every action that is interim rationalizable for a finite type *t* remains interim ε -rationalizable for all types sufficiently close to *t* in the uniform-weak topology. Conversely, given any finite type *t* there exist $\varepsilon > 0$ and a finite game such that some interim rationalizable action for *t* fails to be interim ε -rationalizable for every type that is not close to *t* in the uniform-weak topology. Our results thus establish the equivalence between the uniform-weak topology and the strategic topology of Dekel, Fudenberg, and Morris (2006) around finite types.

1 Introduction

Incomplete information games are games in which some payoff-relevant states are not common knowledge among the players. Harsanyi (1967-68) observes that the Bayesian analysis of incomplete information games requires a model in which each player is

^{*}We thank Jeffrey Ely, Drew Fudenberg, George J. Mailath, Marcin Peski and Qingmin Liu for helpful comments.

equipped with an infinite hierarchy of beliefs: a belief about the payoff-relevant states, a belief about his opponents' beliefs about the payoff-relevant states, and so on. Following this observation, Harsanyi (1967-68) introduces type spaces as a parsimonious model that encodes the belief hierarchies and is suitable for game theoretic analysis, in that interim best-reply sets can be appropriately defined. Mertens and Zamir (1985) provide a foundation for the use of type spaces showing that the space \mathcal{T} of coherent belief hierarchies is a *universal type space*. That is, \mathcal{T} is a type space itself and, moreover, every type space can be embedded in \mathcal{T} via a belief-preserving morphism. Hence, the universal type space \mathcal{T} captures the richness of any abstract type space, and not more.

The Mertens-Zamir universal type space comes with a natural topology: the product topology.¹ A distinctive feature of the product topology is that it is insensitive to the tails of belief hierarchies: two types are close in the product topology if, and only if, their k^{th} -order beliefs are close for some large *finite* $k \ge 1$. Strategic behavior, however, can be very sensitive to high order beliefs. This is true even for interim rationalizability (see Dekel, Fudenberg, and Morris (2007)), the most permissive solution concept consistent with common knowledge of rationality. In effect, in Rubinstein (1989)'s electronic mail game, an action – "attack" – is *strictly* rationalizable for a type *t*, but fails to be rationalizable for all types in a sequence that converges to *t* in the product topology. Hence, to the extent that strategic behavior is what one ultimately cares about, the product topology yields an inadequate notion of proximity of types.

From this point of view, the appropriateness of a topology on types depends on what is meant by strategic behavior. But given a solution concept, it is natural to consider the coarsest topology under which the correspondence that maps types into solutions is continuous in every game. For the solution concept of interim ε -rationalizability this yields the *strategic topology* on types introduced by Dekel, Fudenberg, and Morris (2006), hereafter DFM. The strategic topology, while being strong enough to render ε rationalizable behavior continuous, is remarkably weak: DFM show that finite types are dense.

Given the importance of the strategic topology,² and the fact that it is a topology on types that is independent of the strategic situation (i.e., action sets and payoffs), we find it conceptually important to give it a characterization in terms of properties of the belief hierarchies, with no direct reference to such concepts as behavior strategies

¹It is only when \mathcal{T} is equipped with the measurable structure induced by the product topology that \mathcal{T} can be shown to be a universal type space. This is the sense in which the product topology is natural.

²One reason why the study of strategic convergence seems important is that it appears to be a useful step for the examination of robustness questions in mechanism design.

and best-replies, which are tied to fixed games. In this paper, we take a first step towards such characterization and show that around *finite types* the strategic topology coincides with the *uniform-weak topology*. The latter is the topology induced by the metric *d*, defined as follows: for each order $k \ge 1$, let d^k be a metric that induces the topology of weak convergence of *k*-order beliefs; metric *d* is defined as the supremum of d^k over all orders $k \ge 1$.

The connection between uniform topologies on types and the strategic topology was first suggested by Morris (2002), who studies a particular class of infinite-action games, called higher-order expectation games (HOE), and shows that a certain topology on types (different from ours) is equivalent to the weakest topology under which the ε -rationalizability correspondence is continuous in every game of the HOE class. This uniform topology is too strong for our purposes: there exists a sequence of types, (t_n) , which fails to converge (in this uniform topology) to a type t, and yet in every *finite* game, every rationalizable action for t remains ε -rationalizable for t_n for all n large enough. Hence, the strategic separation of types that are not close in this uniform topology requires an infinite game.

The connection between uniform and strategic convergence of types also underlies the main result in Monderer and Samet (1989). They show that a sufficient condition for the correspondence of Bayesian-Nash ε -equilibrium to be continuous at a *completeinformation* type profile is that the sequence of approximating type profiles converges to its complete information limit in the *common* p-belief sense. (That is, for every p > 0, at every type profile sufficiently far in the tail of the sequence there is common p-belief of the state that is common certainty in the limit.) Moreover, they show that this notion of convergence of type profiles yields strategic continuity in every game. Kajii and Morris (1997) prove the converse: If a sequence of type profiles fails to converge to a complete information type in the common p-belief sense, then a finite game exists such that for some $\varepsilon > 0$, some equilibrium of the complete information game will fail to be an ε -equilibrium at every type profile in the tail of the sequence. It is interesting to note that a sequence of types converges to a complete information type in the uniform-weak topology if, and only if, it converges in the common p-belief sense. Hence, the topology of uniform-weak convergence extends the notion of common p-belief convergence to incomplete information limit types.

This paper is also closely related to contemporaneous work by Ely and Peski (2007). Following their terminology, a type is called *regular* if for every finite game the ε -rationalizability correspondence is continuous in the product topology. Ely and Peski (2007) provide an insightful characterization of regular types in terms of properties of the belief hierarchies and show that the set of regular types is generic (in a topological

sense). They prove:

Theorem (Ely and Peski (2007)). A type t is regular if, and only if, for every p > 0 and every closed (in the product topology), proper subset W of the universal type space, W is not common p-belief at t. Furthermore, the set of regular types is residual, that is, it contains a countable intersection of open and dense sets.

Thus, in a topological sense, around almost all types the strategic topology coincides with the product topology. While topological genericity is interesting, we think it should not be the end of the story. We find it conceptually important to characterize the strategic topology around *critical types*, namely, those types which are not regular. In fact, given Ely and Peski (2007)'s result, it appears to us that every type space ever considered in applications consists entirely of critical types. We take a first step towards such characterization by proving the equivalence between the strategic topology and the uniform-weak topology around *finite types*. All finite types are critical, but not conversely. ³

2 Preliminaries

Hereafter we fix a two-player set I and a finite space of basic uncertainty Θ . Given a player $i \in I$, let -i denote the other player in I. Given a topological space X, write $\Delta(X)$ for the set of probability measures on the Borel subsets of X endowed with the topology of weak convergence of probability measures. Unless explicitly noted, all product spaces will be endowed with the product topology and subspaces with the relative topology.

2.1 The Mertens-Zamir Universal Type Space

Let $Y^0 = \Theta$ and $Y^1 = Y^0 \times \Delta(Y^0)$. Then, for $k \ge 2$, define recursively

$$Y^{k} = \left\{ (\theta, \mu^{1}, \dots, \mu^{k}) \in Y^{0} \times \Delta(Y^{0}) \times \dots \times \Delta(Y^{k-1}) : \operatorname{marg}_{Y^{\ell-2}} \mu^{\ell} = \mu^{\ell-1} \ \forall \ell = 2, \dots, k \right\}$$

By the coherency conditions on marginal distributions from the definition of Y^k , an element of Y^k is uniquely identified by its first and last coordinates. Thus, with slight abuse of notation, given $\theta \in \Theta$ and $\mu^k \in \Delta(Y^{k-1})$, we will sometimes write $(\theta, \mu^k) \in Y^k$.

³We conjecture our characterization is valid for arbitrary critical types, but do not have a proof yet.

The Mertens-Zamir *universal type space* \mathcal{T} is defined as

$$\mathcal{T} = \left\{ (\mu^1, \mu^2, \dots) \in \Delta(Y^0) \times \Delta(Y^1) \times \dots : \operatorname{marg}_{Y^{k-2}} \mu^k = \mu^{k-1} \quad \forall k \ge 2 \right\}.$$

For each $k \ge 1$, let $\pi^k : \mathcal{T} \to \Delta(Y^{k-1})$ denote the natural projection. For every $i \in I$ and $k \ge 1$, let \mathcal{T}_i and Y_i^k denote copies of \mathcal{T} and Y^k , respectively, write $\pi_i^k : \mathcal{T}_i \to \Delta(Y_{-i}^{k-1})$ for π^k , and define $\mathcal{T}_i^k = \pi_i^k(\mathcal{T}_i)$. An element $t_i \in \mathcal{T}_i$ is a *type* of player *i*, and $\pi_i^k(t_i)$ is its associated *k*-order belief.

Each type of *i* uniquely determines a belief over $\Theta \times \mathcal{T}_{-i}$. More precisely, for each $t_i \in \mathcal{T}_i$ there exists a unique probability measure $\mu_i(t_i) \in \Delta(\Theta \times \mathcal{T}_{-i})$ whose marginal on Y_{-i}^{k-1} coincides with $\pi_i^k(t)$ for all $k \ge 1$. Conversely, for every such probability measure in $\Delta(\Theta \times \mathcal{T}_{-i})$ there exists a unique type $t_i \in \mathcal{T}_i$ such that the latter belief-preservation property holds. Moreover, the map $\mu_i : \mathcal{T}_i \to \Delta(\Theta \times \mathcal{T}_{-i})$ is a homeomorphism.

A *finite type space* is a collection $(T_i)_{i \in I}$, with T_i a finite subset of \mathcal{T}_i for all $i \in I$, such that the support of $\mu_i(t_i)$ is contained in $\Theta \times T_{-i}$ for all $t_i \in T_i$ and $i \in I$. A type $t_i \in \mathcal{T}_i$ is called a *finite type* if $t_i \in T_i$ for some finite type space $(T_j)_{j \in I}$.

2.2 Interim Correlated Rationalizability and the Topologies on Types

A finite game is a tuple $(A_i, g_i)_{i \in I}$ with each A_i a finite set and $g_i : A \times \Theta \rightarrow [-1, 1]$, where $A = X_{i \in I}A_i$. For a mixed action profile $\alpha \in \Delta(A_i) \times \Delta(A_{-i})$ write $g(\alpha, \theta)$ for the expectation of g under (α, θ) .

Given a finite game $G = (A_i, g_i)_{i \in I}$ and a type $t_i \in \mathcal{T}_i$, for each $k \ge 0$ we denote by $R_i^k(\varepsilon, t_i; G)$ the set of *k*-order ε -rationalizable actions of type t_i . These sets are defined recursively as follows (see Dekel, Fudenberg, and Morris (2007)):

$$R_i^0(\varepsilon, t_i; G) = A_i$$

and for $k \ge 1$, action $a_i \in A_i$ belongs to $R_i^k(\varepsilon, t_i; G)$ if there exists a measurable function $\sigma_{-i} : \Theta \times \mathcal{T}_{-i} \to \Delta(A_{-i})$ such that:

(a) supp
$$\sigma_{-i}(\theta, t_{-i}) \subseteq R^{k-1}_{-i}(\varepsilon, t_{-i}; G)$$
 for $\mu_i(t_i)$ -almost every (θ, t_{-i}) , and

(b) for all
$$a'_i \in A_i$$

$$\int_{\Theta \times \mathcal{T}_{-i}^{k-1}} \left[g_i(a_i, \sigma_{-i}(\theta, t_{-i}), \theta) - g_i(a_i', \sigma_{-i}(\theta, t_{-i}), \theta) \right] \mu_i(t_i)(d\theta, dt_{-i}) \ge -\varepsilon$$

The set of ε -rationalizable actions of type t_i is then defined as

$$R_i(\varepsilon, t_i; G) = \bigcap_{k \ge 1} R_i^k(\varepsilon, t_i; G)$$

Note that the set $R_i^k(\varepsilon, t_i; G)$ only depends on t_i via the *k*-order beliefs $\pi_i^k(t_i)$ of t_i . Thus, with slight abuse of notation and whenever convenient, given any $t_i^k \in \mathcal{T}_i^k$ we will write $R_i^k(\varepsilon, t_i^k; G)$ to indicate the *k*-order ε -rationalizable strategies of any type of *i* with associated *k*-order beliefs t_i^k .

Definition 2.1. The *strategic topology* is the weakest topology on \mathcal{T}_i such that, for every finite game *G*, the correspondence $(\varepsilon, t_i) \mapsto R_i(\varepsilon, t_i; G)$ is continuous.

Dekel, Fudenberg, and Morris (2006) introduce a distance d_i^S on \mathcal{T}_i that metrizes the strategic topology.

Given a metric space (X, ρ) , the *Prohorov distance* between any two $\mu, \mu' \in \Delta(X)$ is

$$\inf \{ \delta > 0 : \mu'(A) \ge \mu(A^{\delta}) - \delta \text{ for every Borel } A \subseteq X \},\$$

where A^{δ} denotes the set of all $x \in X$ such that $\inf_{y \in A} \rho(x, y) < \delta$.

Now let d^0 be the discrete metric on Θ and write d^1 for the Prohorov distance on $\Delta(\Theta)$. Then, recursively for every $k \ge 2$, let d^k be the Prohorov distance on $\Delta(Y^{k-1})$ when Y^{k-1} is given the product metric induced by $d^0, d^1, \ldots, d^{k-1}$.

Definition 2.2. The *uniform-weak topology* on \mathcal{T} is the topology induced by the metric

$$d(t,t') = \sup_{k \ge 1} d^k \big(\pi^k(t), \pi^k(t') \big) \qquad \text{for all } (t,t') \in \mathcal{T}.$$

Two types are close in the uniform-weak topology if and only if they have similar first-order beliefs, attach similar probabilities to the other player having similar first-order beliefs, and so on, where the degree of similarity is uniform over the levels of the belief hierarchy.

Interestingly, if t_i is a *complete information type*, that is, a type at which there is common knowledge of some state θ , then for all $\delta > 0$ and $t'_i \in \mathcal{T}_i$

$$d(t'_i, t_i) < \delta \iff \theta$$
 is common $(1 - \delta)$ -belief at t'_i .

Hence, the uniform-weak topology is an extention of the notion of common p-belief (Monderer and Samet (1989)) to perturbations of incomplete information environments.

3 Equivalence Between the Strategic and the Uniform-weak Topologies on Types

Proposition 3.1. Around finite types, the uniform-weak topology is stronger than the strategic topology. More precisely, for every player $i \in I$, finite type $t_i \in \mathcal{T}_i$ and $\varepsilon > 0$ there exists $\delta > 0$ such that for all $t'_i \in \mathcal{T}_i$,

$$d_i(t_i, t'_i) < \delta \implies d_i^S(t_i, t'_i) < \varepsilon.$$

The proposition is a direct implication of the following:

Lemma 3.1. Let $G = (A_i, g_i)_{i \in I}$ be a finite game and $(T_i)_{i \in I}$ a finite type space. For every $\varepsilon > 0$ there exists $\delta > 0$ such that for every $k \ge 1$, $i \in I$ and $(t_i, t'_i) \in T_i \times T_i$,

$$d_i^k(\pi_i^k(t_i), \pi_i^k(t_i')) < \delta \implies R_i^k(0, t_i; G) \subseteq R_i^k(\varepsilon, t_i'; G).$$

Proof. For each $i \in I$ and $k \ge 1$, let $T_i^k \equiv \pi_i^k(T_i)$, $\eta_i^k \equiv \min \{d_i^k(t_i^k, T_i^k \setminus \{t_i^k\}) \mid t_i^k \in T_i^k\}$ and $\eta \equiv \min_{i \in I} \min_{k \ge 1} \eta_i^k$. Since $(T_i)_{i \in I}$ is a finite type space, $\eta_i^k > 0$ for all $i \in I$ and $k \ge 1$. Moreover, there exists $k_0 \ge 1$ such that $\eta_i^k = \eta_i^{k_0}$ for all $k \ge k_0$, and hence we have $\eta > 0$. Choose any $0 < \delta < \frac{1}{2} \min \{\eta, \varepsilon |\Theta|^{-1}(|2^{A_i}| + |2^{A_j}|)^{-1}\}$.

The proof proceeds by induction in *k*. Fix $t_i \in T_i$ and $t'_i \in \mathcal{T}_i$ with $d^1_i(\pi^1_i(t_i), \pi^1_i(t_i)) < \delta$. For each $a_i \in R^1_i(0, t_i; G)$ there exists a behavior strategy $b^1_i : \Theta \to \Delta(A_j)$ such that

$$\sum_{\theta\in\Theta}\ell_i^1(\theta;a_i,a_i',b_i^1)\,\pi_i^1(t_i)[\theta] \ge 0,$$

where

-

$$\mathcal{B}_i^0(\theta;a_i,a_i',b_i^1) \equiv \sum_{a_j \in A_j} \left(g_i(a_i,a_j,\theta) - g_i(a_i',a_j,\theta)\right) b_i^1(\theta)[a_j].$$

Then

$$\begin{split} \sum_{\theta \in \Theta} \ell_i^1(\theta; a_i, a_i', b_i^1) \, \pi_i^1(t_i')[\theta] & \geq \sum_{\theta \in \Theta} \ell_i^1(\theta; a_i, a_i', b_i^1) \, (\pi_i^1(t_i')[\theta] - \pi_i^1(t_i)[\theta]) \\ & \geq -\sum_{\theta \in \Theta} \left| \pi_i^1(t_i')(\theta) - \pi_i^1(t_i)(\theta) \right| > -\delta |\Theta| > -\varepsilon \,, \end{split}$$

where the second follows from $|g_i| \leq 1$ and the penultimate inequality follows from $d_i^1(\pi_i^1(t_i), \pi_i^1(t'_i)) < \delta$. Hence $a_i \in R_i^1(\varepsilon, t'_i; G)$, which proves our claim for k = 1.

Now let $k \ge 2$ and assume the claim holds true for k - 1. Let $t_i \in T_i$ and $t'_i \in \mathcal{T}_i$ be a pair of types with $d_i^k(\pi_i^k(t_i), \pi_i^k(t'_i)) < \delta$. Fix an action $a_i \in R_i^k(0, t_i; G)$ and let $b_i^k : \Theta \times T_j^{k-1} \to \Delta(A_j)$ be a *k*-order behavior strategy for type t_i such that:⁴

· for all $(\theta, t_j^{k-1}) \in \Theta \times T_j^{k-1}$,

(

$$b_i^k(\theta, t_j^{k-1}) \in \Delta(R_j^{k-1}(0, t_j^{k-1})),$$
(3.1)

• for all $a'_i \in A_i$,

$$\sum_{\theta, t_j^{k-1}) \in \Theta \times T_j^{k-1}} \ell_i^k(\theta, t_j^{k-1}; a_i, a_i', b_i^k) \, \pi_i^k(t_i) [\theta, t_j^{k-1}] \ge 0, \tag{3.2}$$

where $\ell_i^k(\theta, t_j^{k-1}; a_i, a'_i, b_i^k)$ is the expected payoff loss (under b_i^k) of the deviation from a_i to a'_i conditional on (θ, t_i^{k-1}) . That is,

$$\ell_i^k(\theta, t_j^{k-1}; a_i, a_i', b_i^k) = \sum_{a_j \in A_j} \left(g_i(a_i, a_j, \theta) - g_i(a_i', a_j, \theta) \right) b_i^k(\theta, t_j^{k-1})[a_j]$$

For each $A'_i \subseteq A_j$, let

$$T_j^{k-1}(A_j') \equiv \{t_j^{k-1} \in T_j^{k-1} : R_j^{k-1}(0, t_j^{k-1}) = A_j'\},\$$

so that $\{T_j^{k-1}(A'_j) : A'_j \subseteq A_j, T_j^{k-1}(A'_j) \neq \emptyset\}$ is a partition of T_j^{k-1} . For each $C \subseteq \mathcal{T}_j^{k-1}$ write $B_j^{k-1,\delta}(C)$ for the δ -open-ball around C in $(\mathcal{T}_j^{k-1}, d_j^{k-1})$ (with the convention that $B_j^{k-1,\delta}(\emptyset) = \emptyset$). Since $\delta < \eta/2$, we have $B_j^{k-1,\delta}(T_j^{k-1}(A'_j)) \cap B_j^{k-1,\delta}(T_j^{k-1}(A''_j)) \neq \emptyset$ for every $A'_j, A''_j \subseteq A_j, A'_j \neq A''_j$, such that $T_j^{k-1}(A'_j) \neq \emptyset$ and $T_j^{k-1}(A''_j) \neq \emptyset$.

Consider the *k*-order behavior strategy $b'_i^k : \Theta \times \mathcal{T}_j^{k-1} \to \Delta(A_j)$ for type t'_i defined as follows:

$$\text{If } t_j^{k-1} \in B_j^{k-1,\delta}(T_j^{k-1}(A_j')) \text{ for some } A_j' \subseteq A_j, \text{ then for each } \theta \in \Theta \text{ set}$$
$$b_i'^k(\theta, t_j^{k-1})[\cdot] \equiv \sum_{\tilde{t}_j^{k-1} \in T_j^{k-1}(A_j')} b_i^k(\theta, \tilde{t}_j^{k-1})[\cdot] \pi_i^k(t_i)(\tilde{t}_j^{k-1} | \theta, T_j^{k-1}(A_j')),$$

where $\pi_i^k(t_i)(\cdot | \theta, T_j^{k-1}(A'_j))$ is the conditional probability of $\pi_i^k(t_i)$ on the event $\{\theta\} \times \{T_j^{k-1}(A'_j)\}$.

⁴Since $\Theta \times T_j^{k-1}$ is a finite set, The requirement that (a) hold for all $(\theta, t_j^{k-1}) \in \Theta \times T_j^{k-1}$, rather than only for $(\theta, t_j^{k-1}) \in \Theta \times T_j^{k-1} \in \text{supp } \pi^k(t_i)$, is without loss of generality.

• If $t_j^{k-1} \in \mathcal{T}_j^{k-1} \setminus \bigcup \{B_j^{k-1,\delta}(T_j^{k-1}(A_j')) : A_j' \subseteq A_j\}$, then for each $\theta \in \Theta$ set $b_i'^k(\theta, t_j^{k-1})$ equal to an arbitrary measurable selection from the set-valued map $\tilde{t}_j^{k-1} \mapsto \Delta(R_j^{k-1}(\varepsilon, \tilde{t}_j^{k-1}))$, where the choice of the selection is immaterial for the ensuing argument.⁵

Since the nonempty $B_j^{k-1,\delta}(T_j^{k-1}(A'_j))$'s are pairwise disjoint, b'_i^k is well defined. For all $\theta \in \Theta$ and $t_i^{k-1} \in \mathcal{T}_i^{k-1}$, we claim:

$$b_i^{\prime k}(\theta, t_j^{k-1}) \in \Delta(R_j^{k-1}(\varepsilon, t_j^{k-1})).$$

For $\theta \in \Theta$ and $t_j^{k-1} \in \mathcal{T}_j^{k-1} \setminus \bigcup \{B_j^{k-1,\delta}(T_j^{k-1}(A'_j)) : A'_j \subseteq A_j\}$, the claim follows from the definition of b'_i^k . Fix $\theta \in \Theta$, $A'_j \subseteq A_j$ (with $T_j^{k-1}(A'_j) \neq \emptyset$) and $t_j^{k-1} \in B_j^{k-1,\delta}(T_j^{k-1}(A'_j))$. By construction,

$$\operatorname{supp} b_i^{\prime k}(\theta, t_j^{k-1}) \subseteq \bigcup_{\tilde{t}_i^{k-1} \in T_i^{k-1}(A_j^{\prime})} \operatorname{supp} b_i^{\prime k}(\theta, \tilde{t}_j^{k-1}) \subseteq A_j^{\prime}.$$

Since $t_j^{k-1} \in B_j^{k-1,\delta}(T_j^{k-1}(A'_j))$, we have $d_j^{k-1}(t_j^{k-1}, \tilde{t}_j^{k-1}) < \delta$ for some $\tilde{t}_j^{k-1} \in T_j^{k-1}(A'_j)$. Hence, by the induction hypothesis,

$$A'_{j} = R_{j}^{k-1}(0, \tilde{t}_{j}^{k-1}) \subseteq R_{j}^{k-1}(\varepsilon, t_{j}^{k-1}),$$

and therefore,

$$\operatorname{supp} b_i^{\prime k}(\theta, t_j^{k-1}) \subseteq R_j^{k-1}(\varepsilon, t_j^{k-1}),$$

which proves our claim.

It remains to show that action a_i is ε -optimal (under b'_i^k) for type t'_i . Fix $a'_i \in A_i$ and abbreviate $\ell'_i^k(\theta, t_j^{k-1}) \equiv \ell_i^k(\theta, t_j^{k-1}; a_i, a'_i, b'_i^k)$.

Since
$$d_i^k(\pi^k(t_i), \pi^k(t_i')) < \delta$$
 and supp $\pi^k(t_i) \subseteq T_j^{k-1}$, we have

$$\pi^k(t_i') \Big[\Theta \times \bigcup_{A_i' \subseteq A_i} B_j^{k-1,\delta}(T_j^{k-1}(A_j')) \Big] \ge 1 - \delta > 1 - \varepsilon/2$$

and therefore,

$$\begin{split} \int_{\Theta \times \mathcal{T}_j^{k-1}} \ell'_i^k(\theta, t_j^{k-1}) \ \pi_i^k(t_i') [d\theta, dt_j^{k-1}] \ \geq \\ & \geq \sum_{\theta \in \Theta, A_j' \subseteq A_j} \int_{B_j^{k-1,\delta}(T_j^{k-1}(A_j'))} \ell'_i^k(\theta, t_j^{k-1}) \ \pi_i^k(t_i') \Big[\{\theta\}, dt_j^{k-1}\Big] - \frac{\varepsilon}{2}. \end{split}$$

⁵Since $\Delta(R_j^{k-1}(\varepsilon, \cdot))$ is upper hemi-continuous, the existence of a measurable selection follows from the Kuratowski-Ryll-Nardzewski Theorem (see Aliprantis and Border (1999)).

But since $b'_i{}^k(\theta, \cdot)$ is constant on each $B_j^{k-1,\delta}(T_j^{k-1}(A'_j))$ for every fixed $\theta \in \Theta$, we have

$$\int_{\Theta \times \mathcal{T}_{j}^{k-1}} \ell'_{i}^{k}(\theta, t_{j}^{k-1}) \pi_{i}^{k}(t_{i}') [d\theta, dt_{j}^{k-1}] \geq$$

$$\geq \sum_{\theta \in \Theta, A_{j}' \in A_{j}} \overline{\ell}_{i}^{k}(\theta, A_{j}') \pi_{i}^{k}(t_{i}') \Big[\{\theta\} \times B_{j}^{k-1,\delta}(T_{j}^{k-1}(A_{j}')) \Big] - \frac{\varepsilon}{2}, \quad (3.3)$$

where $\overline{\ell}_i^k(\theta, A'_j) \equiv \ell'_i^k(\theta, t_j^{k-1})$ for any, and hence all, $t_j^{k-1} \in B_j^{k-1,\delta}(T_j^{k-1}(A'_j))$.

On the other hand, it follows from the definition of b'_i^k , iterated expectations and (3.2) that

$$\begin{split} \sum_{\theta \in \Theta, A'_j \subseteq A_j} \overline{\ell}_i^k(\theta, A'_j) \ \pi_i^k(t_i) \Big[\{\theta\} \times T_j^{k-1}(A'_j) \Big] \ = \\ \sum_{(\theta, t_j^{k-1}) \in \Theta \times T_j^{k-1}} \ell_i^k(\theta, t_j^{k-1}; a_i, a'_i, b^k_i) \ \pi_i^k(t_i) \big[\theta, t_j^{k-1}\big] \ge 0. \end{split}$$

Therefore,

$$\sum_{\theta \in \Theta, A'_{j} \subseteq A_{j}} \overline{\ell}_{i}^{k}(\theta, A'_{j}) \pi_{i}^{k}(t'_{i}) \Big[\{\theta\} \times B_{j}^{k-1,\delta}(T_{j}^{k-1}(A'_{j})) \Big] \geq \sum_{\theta \in \Theta, A'_{j} \subseteq A_{j}} \overline{\ell}_{i}^{k}(\theta, A'_{j}) \Big(\pi_{i}^{k}(t'_{i}) \Big[\{\theta\} \times B_{j}^{k-1,\delta}(T_{j}^{k-1}(A'_{j})) \Big] - \pi_{i}^{k}(t_{i}) \Big[\{\theta\} \times T_{j}^{k-1}(A'_{j}) \Big] \Big) > - |\Theta| |2^{A_{j}} |\delta\rangle - \varepsilon/2, \quad (3.4)$$

where the penultimate inequality follows from $d_i^k(\pi_i^k(t_i), \pi_i^k(t_i')) < \delta$ and $|g_i| \le 1$. Combining (3.3) and (3.4) yields

$$\int_{\Theta \times \mathcal{T}_j^{k-1}} \ell'_i^k(\theta, t_j^{k-1}) \, \pi_i^k(t_i') [d\theta, dt_j^{k-1}] \geq -\varepsilon,$$

and therefore $a_i \in R_i^k(\varepsilon, t'_i; G)$, as required.

Proposition 3.2. Around finite types, the uniform-weak topology is weaker than the strategic topology. More precisely, for every player $i \in I$, finite type $t_i \in \mathcal{T}_i$ and $\delta > 0$ there exists $\varepsilon > 0$ such that for all $t'_i \in \mathcal{T}_i$,

$$d_i(t_i, t_i') > \delta \implies d_i^S(t_i, t_i') > \varepsilon.$$

The proposition is a direct implication of the following lemma, which relies on Lemma A.1 from appendix A.

Lemma 3.2. Let $(T_i)_{i \in I}$ be a finite type space. For every $\delta > 0$ there exist $\varepsilon > 0$ and a finite game G such that for every $i \in I$, $t_i \in T_i$ and $t'_i \in T_i$,

$$d_i(t_i, t'_i) > \delta \implies R_i(0, t_i; G) \notin R_i(\varepsilon, t'_i; G).$$

Proof. Fix $\delta > 0$. For each $i \in I$, let $\mu_i : T_i \to \Delta(\Theta \times T_j)$ be the belief mapping. Define $\zeta = \delta |\Theta|^{-1} (|T_i|^2 + |T_j|^2)^{-1}$. By Lemma A.1, there exist $\varepsilon > 0$ and a game with finite action sets $A_i \supseteq T_i$ such that for every $t_i \in T_i$:

- (i) t_i is a best-reply to belief $\mu_i(t_i)$ (viewed as a probability over $\Theta \times A_j$);
- (ii) for every belief $\mu'_i \in \Delta(\Theta \times A_j)$, t_i is an ε -best-reply to μ'_i only if $||\mu'_i \mu_i(t_i)|| \leq \zeta$, where $|| \cdot ||$ denotes the maximum norm.

We now claim:

Claim. For every $k \ge 1$, $i \in I$, and $(t_i, t'_i) \in T_i \times \mathcal{T}_i$ with $d^k(\pi^k_i(t_i), \pi^k_i(t'_i)) > \delta$,

(a)
$$t_i \in R_i^k(0, t_i; G);$$

(b)
$$t_i \notin R_i^{\kappa}(\varepsilon, t_i'; G)$$
.

We shall prove the claim by induction in *k*. Consider k = 1 and fix $i \in I$ and $t_i \in T_i$. Let $b_i^1 : \Theta \to \Delta(A_{-i})$ and $v_i \in \Delta(\Theta \times A_{-i})$ be defined as follows: for all $\theta \in \Theta$ and $a_{-i} \in A_{-i}$,

$$b_i^1(\theta)[a_{-i}] = \pi_i^1(t_i)[a_{-i}|\theta]$$
 and $v_i[\theta, a_{-i}] = b_i^1(\theta)[a_{-i}]\pi^1(t_i)[\theta]$.

Since $\pi_i^1(t_i) = \text{marg}_{\Theta} \mu_i(t_i)$, it is clear that $\nu_i = \mu_i(t_i)$. But then it follows from (i) that $t_i \in R_i^1(0, t_i; G)$, which proves part (a) of the claim for k = 1.

To prove part (b), fix an arbitrary $t'_i \in \mathcal{T}_i$ and assume $d^1((\pi^1_i(t_i), \pi^1_i(t'_i)) > \delta$. The latter means $\pi^1_i(t'_i)[\Theta'] < \pi^1_i(t_i)[\Theta'] - \delta$ for some $\Theta' \subseteq \Theta$, hence

$$\pi_i^1(t_i')[\theta] < \pi_i^1(t_i)[\theta] - \frac{\delta}{|\Theta|} \quad \text{for some } \theta \in \Theta.$$
(3.5)

Now fix any $b_i^1 : \Theta \to \Delta(A_{-i})$ and define $v_i \in \Delta(\Theta \times A_{-i})$ as follows:

$$v_i[\theta, a_{-i}] = b_i^1(\theta)[a_{-i}] \pi^1(t'_i)[\theta]$$
 for all $\theta \in \Theta$ and $a_{-i} \in A_{-i}$.

Pick $\theta \in \Theta$ satisfying (3.5). Then, since marg_{Θ} $\nu_i = \pi_i^1(t'_i)$ and marg_{Θ} $\mu_i(t_i) = \pi_i^1(t_i)$,

$$\sum_{a_{-i}\in T_{-i}} v_i[\theta, a_{-i}] \leq \sum_{a_{-i}\in A_{-i}} v_i[\theta, a_{-i}] \qquad (\text{since } T_{-i} \subseteq A_{-i})$$
$$< \sum_{a_{-i}\in A_{-i}} \mu_i(t_i)[\theta, a_{-i}] - \frac{\delta}{|\Theta|} \qquad (\text{by (3.5)})$$
$$= \sum_{a_{-i}\in T_{-i}} \mu_i(t_i)[\theta, a_{-i}] - \frac{\delta}{|\Theta|} \qquad (\text{as } \mu_i(t_i)[\{\theta\} \times (A_{-i} \setminus T_{-i})] = 0)$$

But then $v_i[\theta, a_{-i}] < \mu_i(t_i)[\theta, a_{-i}] - \delta(|\Theta||T_{-i}|)^{-1}$ for some $a_{-i} \in T_{-i}$, hence

$$||\mu_i(t_i) - \nu_i|| > \frac{\delta}{|\Theta||T_{-i}|} > \zeta$$

and, using (ii), also $t_i \notin R_i^1(\varepsilon, t'_i; G)$. This concludes the proof of the claim for k = 1.

Now let $k \ge 2$ and suppose the claim holds true for k-1. Fix $i \in I$ and $t_i \in T_i$. Define the conjecture $\sigma_{-i} : \Theta \times \mathcal{T}_{-i} \to \Delta(A_{-i})$ as follows: for all $(\theta, t_{-i}^{k-1}) \in \Theta \times \mathcal{T}_{-i}^{k-1}$,

$$b_i^k(\theta, t_j^{k-1})[a_j] = \frac{\mu_i(t_i)[\{\theta\} \times (\{a_j\} \cap (\pi_j^{k-1})^{-1}(t_j^{k-1}))]}{\mu_i(t_i)[\{\theta\} \times (\pi_j^{k-1})^{-1}(t_j^{k-1})]} \quad \text{for all } a_j \in T_j$$

if $\mu_i(t_i)[\{\theta\} \times (\pi_j^{k-1})^{-1}(t_j^{k-1})] > 0$, and

$$b_i^k(\theta, t_j^{k-1})[a_j] = 1 / |R_j^{k-1}(0, t_j; G)|$$
 for all $a_j \in R_j^{k-1}(0, t_j; G)$

otherwise. Note that $b_i^k(\theta, t_j^{k-1})[a_j] > 0$ only if $a_j = \hat{t}_j$ for some $\hat{t}_j \in T_j$ with $\pi^{k-1}(\hat{t}_j) = \pi^{k-1}(t_j)$. Therefore, by the induction hypothesis,

$$\operatorname{supp} b_i^k(\theta,t_j^{k-1}) \subseteq R_j^{k-1}(0,t_j^{k-1};G)$$

for all $(\theta, t_j^{k-1}) \in \Theta \times T_j^{k-1}$.

Next, define $\hat{\mu}_i \in \Delta(\Theta \times A_j)$ as

$$\hat{\mu}_i[\theta, a_j] = b_i^1(\theta)[a_j] \pi^1(t_i)[\theta] \quad \text{ for all } \theta \in \Theta \text{ and } a_j \in A_j.$$

Consider the behavior strategy $b_i^k : \Theta \times T_j^{k-1} \to \Delta(A_j)$ defined by:

$$b_i^k(\theta, t_j^{k-1})[a_j] = \mu_i(t_i)[a_j \mid \theta, t_j^{k-1}],$$

for all $(\theta, t_j^{k-1}) \in \Theta \times T_j^{k-1}$ and $a_j \in A_j$.

Behavior strategy b_i^k together with *k*-order beliefs $\pi_i^k(t_i) \in \Delta(\Theta \times T_j^{k-1})$ induce a probability $\hat{\mu}_i \in \Delta(\Theta \times A_j)$ via:

$$\hat{\mu}_{i}[\theta, a_{j}] = \sum_{t_{j}^{k-1} \in T_{j}^{k-1}} b_{i}^{k}(\theta, t_{j}^{k-1})[a_{j}] \pi^{k}(t_{i})[\theta, t_{j}^{k-1}],$$

for all $\theta \in \Theta$ and $a_j \in A_j$. Since $\pi^k(t_i) = \operatorname{marg}_{\Theta \times T_i^{k-1}} \mu_i(t_i)$, we have

$$\hat{\mu}_i[\theta, a_j] = \sum_{\substack{t_j^{k-1} \in T_j^{k-1} \\ = \mu_i(t_i)[\theta, a_j],}} \mu_i(t_i)[a_j | \theta, t_j^{k-1}] \left(\operatorname{marg}_{\Theta \times T_j^{k-1}} \mu_i(t_i) \right) [\theta, t_j^{k-1}]$$

By (i) we have $t_i \in R_i^k(0, t_i; G)$, which proves part (a) of the claim.

Consider part (b). Fix $i \in I$ and $t'_i \in \mathcal{T}_i$ with $d^k(\pi^k_i(t_i), \pi^k_i(t'_i)) > \delta$. Let $b^k_i : \Theta \times \mathcal{T}_j^{k-1} \to \Delta(A_j)$ be an arbitrary behavior strategy such that:

$$b_i^k(\theta, \hat{t}_j^{k-1}) \in \Delta(R_j^{k-1}(\varepsilon, \hat{t}_j^{k-1}; G))$$
(3.6)

for all $(\theta, \hat{t}_j^{k-1}) \in \Theta \times \mathcal{T}_j^{k-1}$. Note that, by the induction hypothesis, for every $t_j \in T_j$ we can have $b_i^k(\theta, \hat{t}_j^{k-1})[t_j] > 0$ only if $d^{k-1}(\hat{t}_j^{k-1}, \pi_j^{k-1}(t_j)) \leq \delta$.

Behavior strategy b_i^k together with *k*-order beliefs $\pi_i^k(t_i')$ induce a probability $\mu_i' \in \Delta(\Theta \times A_j)$ via:

$$\mu_{i}'[\theta, a_{j}] = \int_{\mathcal{T}_{j}^{k-1}} b_{i}^{k}(\theta, \hat{t}_{j}^{k-1})[a_{j}] \, \pi^{k}(t_{i}')[\theta, d\hat{t}_{j}^{k-1}], \qquad (3.7)$$

for all $\theta \in \Theta$ and $a_j \in A_j$.

Since
$$d^{k}(\pi_{i}^{k}(t_{i}), \pi_{i}^{k}(t_{i}')) > \delta$$
, there exists some $(\overline{\theta}, \overline{t}_{j}^{k-1}) \in \Theta \times T_{j}^{k-1}$ such that
 $\pi_{i}^{k}(t_{i}') \Big[\{\overline{\theta}\} \times \overline{B}_{j}^{k-1,\delta}(\overline{t}_{j}^{k-1}) \Big] < \pi_{i}^{k}(t_{i}) [\overline{\theta}, \overline{t}_{j}^{k-1}] - \frac{\delta}{|\Theta||T_{j}|}.$
(3.8)

Let \overline{t}_j be an arbitrary type in T_j with $\pi_j^{k-1}(\overline{t}_j) = \overline{t}_j^{k-1}$. By the induction hypothesis and (3.6), for every $\hat{t}_j^{k-1} \in \mathcal{T}_j^{k-1}$ we can have $b_i^k(\theta, \hat{t}_j^{k-1})[\overline{t}_j] > 0$ only if $\hat{t}_j^{k-1} \in \overline{B}_j^{k-1,\delta}(\overline{t}_j^{k-1})$. Thus, by (3.7),

$$\mu_i'[\overline{\theta},\overline{t}_j] = \int_{\overline{B}_j^{k-1,\delta}(\overline{t}_j^{k-1})} b_i^k(\overline{\theta},\widehat{t}_j^{k-1})[\overline{t}_j] \pi^k(t_i')[\overline{\theta},d\widehat{t}_j^{k-1}],$$

hence

$$\begin{split} \sum_{\left\{\overline{t}_{j}\in T_{j}: \pi_{j}^{k-1}\left(\overline{t}_{j}\right)=\overline{t}_{j}^{k-1}\right\}} \mu_{i}'[\overline{\theta},\overline{t}_{j}] &\leq \pi^{k}(t_{i}') \Big[\left\{\overline{\theta}\right\} \times \overline{B}_{j}^{k-1,\delta}(\overline{t}_{j}^{k-1})\Big] \\ &< \pi_{i}^{k}(t_{i}) \big[\overline{\theta},\overline{t}_{j}^{k-1}\big] - \frac{\delta}{|\Theta||T_{j}|}, \end{split}$$

where the last inequality is just (3.8). But since

$$\pi_i^k(t_i)[\overline{\theta},\overline{t}_j^{k-1}] = \sum_{\{\overline{t}_j \in T_j : \pi_j^{k-1}(\overline{t}_j) = \overline{t}_j^{k-1}\}} \mu_i(t_i)[\overline{\theta},\overline{t}_j],$$

we have

$$\sum_{\overline{t}_j\in T_j} \left| \mu_i'[\overline{\theta},\overline{t}_j] - \mu_i(t_i)[\overline{\theta},\overline{t}_j] \right| > \frac{\delta}{|\Theta||T_j|},$$

and therefore

$$||\mu_i'-\mu_i(t_i)||>\frac{\delta}{|\Theta||T_j|^2}>\zeta.$$

It follows from (ii) that $t_i \notin R_i^k(\varepsilon, t'_i; G)$. This concludes the proof of the claim.

A Appendix for Section 3

Lemma A.1. For each $i \in I$, let T_i be a finite set and $\mu_i : T_i \to \Delta(\Theta \times T_j)$ a function. For every $0 < \zeta < 1$ there exist $\varepsilon > 0$ and a game with finite action sets $A_i \supseteq T_i$ such that for every $t_i \in T_i$:

- (i) t_i is a best-reply to belief $\mu_i(t_i)$ (viewed as a probability over $\Theta \times A_j$);
- (ii) For every belief $\mu'_i \in \Delta(\Theta \times A_j)$, t_i is an ε -best-reply to μ'_i only if $||\mu'_i \mu_i(t_i)|| \leq \zeta$, where $|| \cdot ||$ denotes the maximum norm.

Proof. Fix $\zeta \in (0, 1)$. Let $f_i : \Theta \times T_j \times \Delta(\Theta \times T_j) \longrightarrow \mathbb{R}$ denote the function defined by

$$f_i(\theta, t_j; \mu') = 2\mu'(\theta, t_j) - \sum_{(\theta', t'_j) \in \Theta \times T_j} (\mu'(\theta', t'_j))^2,$$

for all $(\theta, t_j, \mu') \in \Theta \times T_j \times \Delta(\Theta \times T_j)$, and let $F_i : \Delta(\Theta \times T_j) \times \Delta(\Theta \times T_j) \longrightarrow \mathbb{R}$ be the function defined by

$$F_i(\mu^{\prime\prime},\mu^\prime) = \sum_{(\theta,t_j)\in\Theta\times T_j} f_i(\theta,t_j;\mu^{\prime\prime})\,\mu^\prime(\theta,t_j),$$

for all $(\mu'', \mu') \in \Delta(\Theta \times T_j) \times \Delta(\Theta \times T_j)$.

Let $\eta = \frac{1}{2} \min \left\{ F_i(\mu', \mu') - F_i(\mu'', \mu') : (\mu'', \mu') \in \left(\Delta(\Theta \times T_j) \right)^2, ||\mu' - \mu''|| \ge \frac{\zeta}{2} \right\}$. We have $\eta > 0$, for F_i is continuous and $\mu'' = \mu'$ is the unique maximizer of $F_i(\cdot, \mu')$ on $\Delta(\Theta \times T_j)$ for all μ' . By the uniform continuity of F_i , there exists $\gamma > 0$ such that for all $(\mu'', \mu') \in (\Delta(\Theta \times T_j))^2$,

$$||\mu^{\prime\prime}-\mu^{\prime}|| < \gamma \implies F_i(\mu^{\prime},\mu^{\prime}) - F_i(\mu^{\prime\prime},\mu^{\prime}) < \eta.$$

The compact set $\Delta(\Theta \times T_j)$ can be covered by a finite union of open balls of radius γ . Choose one point in which of these balls and let $A_i \subset \Delta(\Theta \times T_j)$ denote the finite set of chosen points. Enlarge A_i , if necessary, to ensure $A_i \supseteq T_i$. (We identify each $t_i \in T_i$ with $\mu_i(t_i)$.) Thus, for every $\mu' \in \Delta(\Theta \times A_j)$ there exists $a_i \in A_i$ such that $F_i(\mu', \mu') - F_i(a_i, \mu') < \eta$.

Define the payoff function $g_i : \Theta \times A_i \times A_j \longrightarrow \mathbb{R}$,

$$g_i(\theta, a_i, a_j) = \begin{cases} f_i(\theta, a_j; a_i) &: a_i \in A_i, a_j \in T_j \\ -\frac{4}{\zeta} &: a_i \in T_i, a_j \in A_j \setminus T_j \\ -1 &: a_i \in A_i \setminus T_i, a_j \in A_j \setminus T_j \end{cases}$$

We are now in a position to prove part (i) of the lemma. Suppose player *i*'s belief over $\Theta \times A_j$ is given by $\mu_i(t_i)$, for some $t_i \in T_i$. It follows directly from the definition of g_i and the fact that $\mu_i(t_i)[\Theta \times T_j] = 1$ that each action $a_i \in A_i$ yields player *i* an expected payoff of $F_i(a_i, \mu_i(t_i))$. Since $F_i(\mu_i(t_i), \mu_i(t_i)) \ge F_i(a_i, \mu_i(t_i))$ for all $a_i \in A_i$, we conclude that t_i is a best-reply to $\mu_i(t_i)$. This proves part (i).

Fix any $0 < \varepsilon < \min\{\eta(1 - \frac{\zeta}{2}), \frac{\zeta}{2}\}$. We shall prove part (ii) now. Fix $t_i \in T_i$ and $\mu' \in \Delta(\Theta \times A_j)$ with $||\mu' - \mu_i(t_i)|| > \zeta$. Suppose $\mu'(\Theta \times T_j) < 1 - \frac{\zeta}{2}$. (The complementary case will be handled in the next paragraph.) Consider a deviation from t_i to an arbitrary action $a_i \in A_i \setminus T_i$. Since F_i maps into [-1, 1], the gain from this deviation is bounded below by

$$(1-\frac{\zeta}{2})(-2)+\frac{\zeta}{2}(-1+\frac{4}{\zeta})=\frac{\zeta}{2}>\varepsilon,$$

and therefore t_i is not an ε -best-reply to μ' , which concludes the proof of part (ii) in the case $\mu'(\Theta \times T_j) < 1 - \frac{\zeta}{2}$.

Now suppose $\mu'(\Theta \times T_j) \ge 1 - \frac{\zeta}{2}$. Since $||\mu' - \mu_i(t_i)|| > \zeta$, there exists $(\theta, t_j) \in \Theta \times T_j$ such that

$$\left|\mu'[\theta,t_j] - \mu_i(t_i)[\theta,t_j]\right| > \zeta.$$
(A.1)

Consider the conditional probability $\bar{\mu}(\cdot) \equiv \mu'(\cdot | \Theta \times T_j)$. We have

$$\bar{\mu}[\theta, t_j] \ge \mu'[\theta, t_j] = \bar{\mu}[\theta, t_j]\mu'(\Theta \times T_j) \ge \bar{\mu}[\theta, t_j] - \frac{\zeta}{2}, \tag{A.2}$$

and therefore

$$\left\| \mu'[\theta,t_j] - \bar{\mu}[\theta,t_j] \right\| < \frac{\zeta}{2}$$

Hence, by (A.2) and (A.1),

$$\begin{split} \left\| \bar{\mu}[\theta, t_j] - \mu_i(t_i)[\theta, t_j] \right\| &\geq \left\| \mu'[\theta, t_j] - \mu_i(t_i)[\theta, t_j] \right\| - \left\| \mu'[\theta, t_j] - \bar{\mu}[\theta, t_j] \right\| \\ &> \frac{\zeta}{2}, \end{split}$$

which implies $F_i(\bar{\mu}, \bar{\mu}) - F_i(t_i, \bar{\mu}) \ge 2\eta$, by the definition of η . Now pick any $a_i \in A_i$ with $||\bar{\mu} - a_i|| < \gamma$, so that $F_i(a_i, \bar{\mu}) - F_i(\bar{\mu}, \bar{\mu}) > -\eta$, and therefore,

$$F_i(a_i,\bar{\mu})-F_i(t_i,\bar{\mu})>\eta.$$

Hence, the payoff gain of the deviation from t_i to a_i is bounded below by

$$\mu'(\Theta \times T_j) \eta \geq (1-\frac{\zeta}{2})\eta > \varepsilon,$$

and therefore t_i is not an ε -best-reply to μ' , as required.

References

- ALIPRANTIS, C. D., AND K. C. BORDER (1999): *Infinite Dimensional Analysis*. Springer, second edn.
- DEKEL, E., D. FUDENBERG, AND S. MORRIS (2006): "Topologies on Types," *Theoretical Economics*, 1, 275–309.
- —— (2007): "Interim Correlated Rationalizability," *Theoretical Economics*, 2, 15–40.

ELY, J. C., AND M. PESKI (2007): "Critical Types," mimeo., Northwestern University.

- HARSANYI, J. (1967-68): "Games with Incomplete Information Played by Bayesian Players, I-III," *Management Science*, 14, 159–182, 320–334, 486–502.
- KAJII, A., AND S. MORRIS (1997): "Refinements and Higher Order Beliefs: A Unified Survey," mimeo., Princeton University.
- MERTENS, J.-F., AND S. ZAMIR (1985): "Formulation of Bayesian Analysis for Games with Incomplete Information," *International Journal of Game Theory*, 14(1), 1–29.
- MONDERER, D., AND D. SAMET (1989): "Approximating Common Knowledge with Common Beliefs," *Games and Economic Behavior*, 1, 170–190.
- MORRIS, S. (2002): "Typical Types," mimeo., Princeton University.
- RUBINSTEIN, A. (1989): "The Electronic Mail Game: A Game with Almost Common Knowledge," *American Economic Review*, 79, 389–391.