



HIER

Harvard Institute of Economic Research

Discussion Paper Number 2093

Topologies on Types

by

Eddie Dekel, Drew Fudenberg
and Stephen Morris

September 2005

HARVARD UNIVERSITY
Cambridge, Massachusetts

This paper can be downloaded without charge from:
<http://post.economics.harvard.edu/hier/2005papers/2005list.html>

The Social Science Research Network Electronic Paper Collection:
<http://ssrn.com/abstract=798944>

Topologies on Types*

Eddie Dekel
Northwestern University
and Tel Aviv University

Drew Fudenberg
Harvard University

Stephen Morris
Princeton University

First Draft: April 2004
This Draft: August 2005

Abstract

We define and analyze "strategic topologies" on types, under which two types are close if their strategic behavior will be similar in all strategic situations. To operationalize this idea, we adopt interim rationalizability as our solution concept, and define a metric topology on types in the Harsanyi-Mertens-Zamir universal type space. This topology is the coarsest metric topology generating upper and lower hemicontinuity of rationalizable outcomes. While upper strategic convergence is equivalent to convergence in the product topology, lower strategic convergence is a strictly stronger requirement, as shown by the electronic mail game. Nonetheless, we show that the set of "finite types" (types describable by finite type spaces) are dense in the lower strategic topology.

JEL classification and keywords: C70, C72, rationalizability, incomplete information, common knowledge, universal type space, strategic topology.

*Previous versions were circulated under the title "Dense Types". We are grateful for comments from seminar participants at Stanford University, the University of Essex, the Cowles Foundation conference on Uncertainty in Economic Theory and the Third International PIER/IGIER Conference on Economic Theory. We are grateful to Dirk Bergemann, Jeff Ely, Yossi Feinberg, Aviad Heifetz, Atsushi Kajii, Bart Lipman, Aldo Rustichini, Muhamet Yildiz and Shmuel Zamir for helpful discussions; Marcin Peski for detailed comments; and Satoru Takahashi for expert proof reading. This material is based upon work supported by the National Science Foundation under Grants No. 0426199 and 0095321.

1 Introduction

The universal type space proposed by Harsanyi (1967/68) and constructed by Mertens and Zamir (1985) shows how (under some technical assumptions) any incomplete information about a strategic situation can be embedded in a single "universal" type space. As a practical matter, applied researchers do not work with that type space but with smaller subsets of the universal type space. Mertens and Zamir showed that finite type spaces are dense under the product topology, but under this topology the rationalizable actions of a given type may be very different from the rationalizable actions of a sequence of types that approximate it.¹ This leads to the question of whether and how one can use smaller type spaces to approximate the predictions that would be obtained from the universal type space.

To address this question, we define and analyze "strategic topologies" on types, under which two types are close if their strategic behavior is similar in all strategic situations. There are three ingredients that need to be formalized in this approach: how we vary the "strategic situations", what "strategic behavior" do we consider (i.e., what solution concept), and what we mean by "similar."

To define "strategic situations," we start with a given space of uncertainty, Θ , and a type space over that space. Holding these fixed, we then consider all possible finite action games where payoffs depend on the actions chosen by the players and the state θ of Nature. Thus we vary the game while holding the type space fixed. This is at odds with a broader interpretation of the universal type space which describes all possible uncertainty—including payoff functions and actions: see the discussion in Mertens, Sorin and Zamir (1994, Remark 4.20b). According to this latter view one cannot identify "higher order beliefs" independent of payoffs in the game. In contrast, our definition of a strategic topology relies crucially on making this distinction. We start with higher-order beliefs about abstract payoff-relevant states Θ , then allow payoffs to vary by changing how payoffs depend on Θ . We are thus implicitly assuming that any "payoff relevant state" can be associated with any payoffs and actions. This is analogous to Savage's assumption that all acts are possible, and thus implicitly that any "outcome" is consistent with any payoff-relevant state. This separation is natural for studying the performance of different mechanisms under various informational

¹This is closely related to the difference between common knowledge and mutual knowledge of order n that is emphasized by Geanakoplos and Polemarchakis (1982) and Rubinstein (1989).

assumptions.²

Our notion of "strategic behavior" is the set of (correlated) interim rationalizable actions that we analyzed in Dekel, Fudenberg and Morris (2005). The rationalizable actions are those obtained by the iterative deletion of all actions that are not best responses given a type's beliefs over others' types and Nature and any (perhaps correlated) beliefs about which actions are played at a given type profile and payoff-relevant state. Under correlated interim rationalizability, a player's beliefs allow for arbitrary correlation between other players' actions and the payoff state; in the complete information case, this definition reduces to the standard definition of correlated rationalizability (e.g., as in Brandenburger and Dekel (1987)).

One might be interested in using equilibrium behavior, as opposed to rationalizability, as the benchmark for rational play. However, the restriction to equilibrium only has bite when additional structure, such as a common prior, is imposed. This is because the set of interim rationalizable actions for a given type t_i is equal to the set of actions that might be played in some equilibrium in a larger type space by a type that has the same beliefs and higher order beliefs about Nature as type t_i .³ Thus, while one might be interested in using equilibrium behavior, as opposed to rationalizability, as the benchmark for rational play, the restriction to equilibrium only has bite when additional structure, such as a common prior, is imposed. Moreover, the set of interim rationalizable actions depends only on types' beliefs and higher-order beliefs about Nature.⁴

It remains to explain our notion of "similar" behavior. Our goal is to find a topology on types that is fine enough that the set of correlated rationalizable actions has the continuity properties that the best response correspondences, rationalizable actions, and Nash equilibria

²Wilson (1987) argued that an important task of mechanism design is to identify mechanisms that perform well under a variety of informational assumptions, and a number of recent works have pursued this agenda, see, e.g., Neeman (2004) and Bergemann and Morris (2004).

³See Dekel, Fudenberg and Morris (2005); this is analogous to the property identified by Brandenburger and Dekel (1987) in the context of complete information.

⁴As noted by Dekel, Fudenberg and Morris (2005) and Ely and Peski (2005), the set of independent interim rationalizable actions can vary across types with the same beliefs and higher order beliefs about Nature. This is because redundant types may differ in the extent they believe that others' actions are correlated with payoff relevant states.

all have in complete-information games, while still being coarse enough to be useful.⁵ A review of those properties helps clarify our work. Fix a family of complete information games with payoff functions u that depend continuously on a parameter λ , where λ lies in a metric space Λ . Because best responses include the case of exact indifference, the set of best responses for player i to a fixed opponents' strategy profile a_j , denoted $BR_i(a_j, \lambda)$, is upper hemi-continuous but not lower hemi-continuous in λ , i.e. it may be that $\lambda^n \rightarrow \lambda$, and $a_i \in BR_i(a_j, \lambda)$, but there is no sequence $a_i^n \in BR_i(a_j, \lambda^n)$ that converges to a_i . However, the set of ε -best responses,⁶ $BR_i(a_j, \varepsilon, \lambda)$, is well behaved: if $\lambda^n \rightarrow \lambda$, and $a_i \in BR_i(a_j, \varepsilon, \lambda)$, then for any $a_i^n \rightarrow a_i$ there is a sequence $\varepsilon^n \rightarrow 0$ such that $a_i^n \in BR_i(a_j, \varepsilon + \varepsilon^n, \lambda^n)$. In particular, the smallest ε for which $a_i \in BR_i(a_j, \varepsilon, \lambda)$ is a continuous function of λ . Moreover the same is true for the set of all ε -Nash equilibria (Fudenberg and Levine (1986)) and for the set of ε -rationalizable actions.⁷ That is, the ε that measures the departure from best response or equilibrium is continuous. Our goal is to find the coarsest topology on types such that the same property holds.

Thus for a fixed game and action, we identify for each type of a player, the smallest ε for which the action is rationalizable. The distance between a pair of types (for a fixed game and action) is the difference between those smallest ε . Our strategic topology requires that, for any convergent sequence, this distance tends to zero pointwise for any action and game. We identify a metric for this topology. The strategic topology requires both a lower hemi-continuity property (the smallest ε does not jump down in limit) and an upper hemi-continuity property (the smallest ε does not jump up in the limit). We show that the upper convergence is equivalent to convergence in the product topology (theorem 2) and that lower convergence implies product convergence, and thus upper convergence (theorem 1).

Our main result is that finite types are dense in the strategic topology (theorem 3). Thus finite type spaces do approximate the universal type space, so that the strategic behavior of

⁵Topology P is finer than topology P' if every open set in P is contained in an open set in P' . The use of a very fine topology such as the discrete topology makes continuity trivial, but it also makes it impossible to approximate one type with another; hence our search for a relatively coarse topology. We will see that our topology is the coarsest metrizable topology with the desired continuity property; this leaves open whether there are other, non-metrizable, topologies that have this property

⁶An action is an ε -best response if it gives a payoff within ε of the best response.

⁷This can be shown by using the equivalence between rationalizable actions and a posteriori equilibria (Brandenburger and Dekel (1987)) and then applying the result for Nash equilibrium.

any type can be approximated by a finite type. However, this does not imply that the set of finite types is large. In fact, while finite types are dense in the strategic topology (and the product topology), they are small in the sense of being category 1 in the product topology and the strategic topology .

Our paper follows Monderer and Samet (1996) and Kajii and Morris (1997) in seeking to characterize "strategic topologies" that yield lower and upper hemi-continuity of strategic outcomes. The earlier papers defined topologies on common prior information systems with a countable number of types, and common p -belief was central to the characterizations.⁸ This paper performs an analogous exercise on the set of all types in the universal type space without the common prior assumption. We do not have a characterization of this strategic topology in terms of beliefs, so we are unable to pin down the relation to these earlier papers.

The paper is organized as follows. Section 2 reviews the electronic mail game and the failure of lower hemicontinuity (but not upper hemicontinuity) of rationalizable outcomes with respect to the product topology. The universal type space is described in section 3 and the incomplete information games and interim rationalizable outcomes we will analyze are introduced in section 4. The strategic topology is defined in section 5 and our main results about the strategic topology are reported in section 6. The concluding section, 7, contains some discussion of the interpretation of our results, the "genericity" of finite types and an alternative stronger uniform strategic topology on types.

2 Electronic Mail Game

To introduce the basic issues we use a variant of Rubinstein's (1989) electronic mail game that illustrates the failure of lower hemi-continuity in the product topology (defined formally below). Specifically, we will use it to provide a sequence of types, t_{ik} , that converge to a type $t_{i\infty}$ in the product topology, while there is an action that is rationalizable for $t_{i\infty}$ but is not ε_k -rationalizable for t_k for any sequence $\varepsilon_k \rightarrow 0$. Thus the set "blows up" at the limit, and the lower-hemicontinuity property discussed in the introduction is not satisfied. Intuitively, for rationalizable play, the tails of higher order beliefs matter, but the product topology is

⁸For Monderer and Samet (1986), an information system was a collection of partitions on a fixed state space with a given prior. For Kajii and Morris (1997), an information system was a prior on a fixed type space.

insensitive to the tails.

On the other hand the set of rationalizable actions does satisfy an upper hemi-continuity property with respect to the sequence of types t_{1k} converging in the product topology to $t_{1\infty}$: fix any sequence $\varepsilon_k \rightarrow 0$ and suppose some action is ε_k -rationalizable for type t_{1k} for all k ; then that action is 0-rationalizable for type $t_{1\infty}$. In section 6 we show that product convergence is equivalent to this upper hemi-continuity property in general.

Example: Each player has two possible actions $A_1 = A_2 = \{N, I\}$ ("not invest" or "invest"). There are two payoff states, $\Theta = \{0, 1\}$. In payoff state 0, payoffs are given by the following matrix:

$\theta = 0$	N	I
N	0, 0	0, -2
I	-2, 0	-2, 2

In payoff state 1, payoffs are given by:

$\theta = 1$	N	I
N	0, 0	0, -2
I	-2, 0	1, 1

Player i 's types are $T_i = \{t_{i1}, t_{i2}, \dots\} \cup \{t_{i\infty}\}$. Beliefs are generated by the following common prior on the type space:

$\theta = 0$	t_{21}	t_{22}	t_{23}	t_{24}	\dots	$t_{2\infty}$
t_{11}	$(1 - \delta) \alpha$	0	0	0	\dots	0
t_{12}	0	0	0	0	\dots	0
t_{13}	0	0	0	0	\dots	0
t_{14}	0	0	0	0	\dots	0
\vdots	\vdots	\vdots	\vdots	\vdots	\ddots	\vdots
$t_{1\infty}$	0	0	0	0	\dots	0

$\theta = 1$	t_{21}	t_{22}	t_{23}	t_{24}	\dots	$t_{2\infty}$
t_{11}	0	0	0	0	\dots	0
t_{12}	$(1 - \delta) \alpha (1 - \alpha)$	$(1 - \delta) \alpha (1 - \alpha)^2$	0	0	\dots	0
t_{13}	0	$(1 - \delta) \alpha (1 - \alpha)^3$	$(1 - \delta) \alpha (1 - \alpha)^4$	0	\dots	0
t_{14}	0	0	$(1 - \delta) \alpha (1 - \alpha)^5$	$(1 - \delta) \alpha (1 - \alpha)^6$	\dots	0
\vdots	\vdots	\vdots	\vdots	\vdots	\ddots	\vdots
$t_{1\infty}$	0	0	0	0	\dots	δ

where $\alpha, \delta \in (0, 1)$. There is an intuitive sense in which the sequence $(t_{1k})_{k=1}^{\infty}$ converges to $t_{1\infty}$. Observe that type t_{12} of player 1 knows that $\theta = 1$ (but does not know if player 2 knows it). Type t_{13} of player 1 knows that $\theta = 1$, knows that player 2 knows it (and knows that 1 knows it), but does not know if 2 knows that 1 knows that 2 knows it. For $k \geq 3$, each type t_{1k} knows that $\theta = 1$, knows that player 2 knows that 1 knows... ($k - 2$ times) that $\theta = 1$. But for type $t_{1\infty}$, there is common knowledge that $\theta = 1$. Thus type t_{1k} agrees with type $t_{1\infty}$ up to $2k - 3$ levels of beliefs. We will later define more generally the idea of *product convergence* of types, i.e., the requirement that k th level beliefs converge for every k . In this example, $(t_{1k})_{k=1}^{\infty}$ converges to $t_{1\infty}$ in the *product topology*.

We are interested in the ε -rationalizable actions in this game. We will provide a formal definition shortly, but the idea is that we will iteratively delete an action for a type at round k if that action is not an ε -best response for any belief over the action-type pairs of the opponent that survived to round $k - 1$.

Clearly, both N and I are 0-rationalizable for types $t_{1\infty}$ and $t_{2\infty}$ of players 1 and 2, respectively. But action N is the unique ε -rationalizable action for all types of each player i except $t_{i\infty}$, for every $\varepsilon < \frac{1+\alpha}{2-\alpha}$ (note that $\frac{1+\alpha}{2-\alpha} > \frac{1}{2}$). Clearly, I is not ε -rationalizable for type t_{11} , since the expected payoff from action N is 0 independent of player 2's action, whereas the payoff from action I is -2 . Now suppose we can establish that I is not ε -rationalizable for types t_{11} through t_{1k} . Type t_{2k} 's expected payoff from action I is at most

$$\frac{1 - \alpha}{2 - \alpha} (1) + \frac{1}{2 - \alpha} (-2) = -\frac{1 + \alpha}{2 - \alpha} < -\frac{1}{2}.$$

Thus I is not ε -rationalizable for type t_{2k} . A symmetric argument establishes if I is not ε -rationalizable for types t_{21} through t_{2k} , then I is not ε -rationalizable for type $t_{1,k+1}$. Thus the conclusion holds by induction. \square

We conclude that strategic outcomes are not continuous in the product topology. Thus

the denseness of finite types in the product topology does not imply that they will be dense in our strategic topology. However, this version of the e-mail game is not a counterexample to the denseness of finite types, since the limit type $t_{1\infty}$ is itself a finite type.

3 Types

In games of incomplete information, a player's chosen action can depend not only on the player's payoff function, but also on his beliefs about the opponents' payoffs, his beliefs about the opponents' beliefs, and so on: in short, on the infinite hierarchy of beliefs. To model this, we look at type spaces that are subsets of the "universal type space" constructed by Mertens and Zamir (1985).⁹

There are two agents, 1 and 2; we denote them by i and $j = 3 - i$.¹⁰ Let Θ be a finite set representing possible payoff-relevant moves by Nature.¹¹ Throughout the paper, we write $\Delta(S)$ for the set of probability measures on the Borel field of any topological space S . Let

$$\begin{aligned} X_0 &= \Theta \\ X_1 &= X_0 \times \Delta(X_0) \\ &\vdots \\ X_k &= X_{k-1} \times \Delta(X_{k-1}) \\ &\vdots \end{aligned}$$

where $\Delta(X_k)$ is endowed with the topology of weak convergence of measures (i.e. the "weak" topology) and each X_k is given the product topology over its two components. A type t is a hierarchy of beliefs $t = (\delta_1, \delta_2, \dots) \in \times_{k=0}^{\infty} \Delta(X_k)$. We are interested in the set of belief

⁹See also Brandenburger and Dekel (1993), Heifetz (1993), and Mertens, Sorin and Zamir (1994). Since Θ is finite, the construction here yields the same universal space with the same σ -fields as the topology-free construction of Heifetz and Samet (1998); see Dekel, Fudenberg and Morris (2005, Lemma 1). Since we are interested in finding a topology, the Heifetz and Samet topology-free approach is more suitable than assuming a topology for this construction. But since the measure-theoretic properties of the two approaches are equivalent we present our work using the more familiar topological constructions.

¹⁰We restrict the analysis to the two player case for notational convenience. We do not think that there would be any difficulties extending the results to any finite number of players.

¹¹We choose to focus on finite Θ as here the choice of a topology is obvious, while in larger spaces the topology on types will depend on the underlying topology on Θ .

hierarchies in which it is common knowledge that the marginal of beliefs δ_k on any level $X_{k'}$, for $k' \leq k - 1$ is equal to $\delta_{k'+1}$. We write T^* for the set of all such types and denote by $\delta_k(t)$ the k th coordinate of $t = (\delta_1, \delta_2, \dots)$. Mertens and Zamir (1985) characterize and prove the existence of induced beliefs for each type over $T^* \times \Theta$, denoted $\pi^* : T^* \rightarrow \Delta(T^* \times \Theta)$; the universal type space is $(T_i^*, \pi_i^*)_{i=1}^2$.

We will focus on types which give rise to distinct hierarchies of beliefs, and ignore the possibility of what Mertens and Zamir (1985) labelled "redundant types"—i.e., types who agree about their higher order beliefs and disagree only in their beliefs over others' redundant types. Previous work (Bergemann and Morris (2001), Battigalli and Siniscalchi (2003), Dekel, Fudenberg, and Morris (2005), Ely and Peski (2005)) has emphasized that redundant types matter for certain solution concepts (such as Nash equilibrium, correlated equilibrium and independent interim rationalizability). In contrast, Dekel, Fudenberg, and Morris (2005) show that two types with the same hierarchy of beliefs, i.e., that map to the same point in the universal type space, have the same set of (correlated) interim rationalizable actions. This is the reason that we use (correlated) interim rationalizability as the solution concept.

Our main result will concern "finite types," meaning types in the universal type space that belong to finite belief-closed subset of the universal type space. To define them more precisely, we need to define *finite* type spaces and discuss how they can be embedded in the universal type space.

Definition 1 *A finite type space is any collection $\mathcal{T} = (T_i, \pi_i)_{i=1}^2$, where each T_i is finite and each $\pi_i : T_i \rightarrow \Delta(T_j \times \Theta)$.*

For any finite belief-closed type space $(T_i, \pi_i)_{i=1}^2$, and any $t_i \in T_i$, we can calculate the infinite hierarchy of i 's beliefs about Θ , beliefs about beliefs, etc. Thus each type t_i can be identified with a unique type $t \in T^*$. Formally, for each $k = 1, 2, \dots$, we can define k th level beliefs iteratively as follows. Let

$$\hat{\pi}_{i1} : T_i \rightarrow \Delta(\Theta)$$

be defined as:

$$\hat{\pi}_{i1}[t_i](\theta) = \sum_{t_j \in T_j} \hat{\pi}_i[t_i](t_j, \theta);$$

Thus $\hat{\pi}_{i1} : T_i \rightarrow \Delta(X_0)$. Now inductively let

$$\hat{\pi}_{ik} : T_i \rightarrow \Delta(X_{k-1})$$

be defined as:

$$\widehat{\pi}_{ik}[t_i](\delta_{j,k-1}, \theta) = \sum_{\{t_j \in T_j : \widehat{\pi}_{j,k-1}(t_j) = \delta_{j,k-1}\}} \widehat{\pi}_i[t_i](t_j, \theta);$$

Thus $\widehat{\pi}_{ik} : T_i \rightarrow \Delta(X_{k-1})$. Let $\widehat{\pi}_i^*(t_i) = (\widehat{\pi}_{ik}(t_i))_{k=1}^\infty$, so $\widehat{\pi}_i^* : T_i \rightarrow T^*$.

We say that a type $t \in T^*$ is "finite" if there exists a finite type space \mathcal{T} containing a type t_i with the same higher-order beliefs as t (i.e., $\widehat{\pi}_i^*(t_i) = t$). That is, a type is finite if it belongs to a finite belief-closed subset of the universal type space.

The most commonly used topology on the universal type space is the "product topology" on the hierarchy:

Definition 2 $t_i^n \rightarrow^* t_i$ if, for each k , $\delta_{ik}(t_i^n) \rightarrow \delta_{ik}(t_i)$ as $n \rightarrow \infty$.

Here, the convergence of beliefs at a fixed level in the hierarchy, represented by \rightarrow , is with respect to the topology of weak convergence of measures. This notion of convergence generates a topology that has been the most used in the study of the universal type space. See, for example, Lipman (2003) and Weinstein and Yildiz (2003). However, our purpose is to construct a topology that reflects strategic behavior, and as the electronic mail game suggests, the product topology is too coarse to have the continuity properties we are looking for.

4 Games and Interim Rationalizability

Our interest is in the implications of these types for play in the family of games \mathcal{G} defined as follows. A game G consists of, for each player i , a finite set of possible actions A_i and a payoff function g_i , where $g_i : A \times \Theta \rightarrow [-M, M]$ and M is an exogenous bound on the scale of the payoffs. Here we restate definitions and results from our companion paper, Dekel, Fudenberg and Morris (2005). In that paper, we varied the type space and held fixed the game G being played and the " ε " in the definition of ε -best response. In this paper, we fix the type space to be the universal type space (and finite belief-closed subsets of it), but we vary the game G and parameter ε , so we make the dependence of the solution on G and ε explicit.

For any subset of actions for all types, we first define the best replies when beliefs over opponents' strategies are restricted to those actions. For any measurable strategy profile

of the opponents, $\sigma_j : T_j^* \times \Theta \rightarrow \Delta(A_j)$, where throughout $j \neq i$, and any belief over opponents' types and the state of Nature, $\pi_i^*(t_i) \in \Delta(T_j^* \times \Theta)$, denote the induced belief over the space of types, Nature and actions by $\nu(\pi_i^*(t_i), \sigma_j) \in \Delta(T_j^* \times \Theta \times A_j)$, where for measurable $F \subset T_j^*$, $\nu(\pi_i^*(t_i), \sigma_j)(F \times \{\theta, a_j\}) = \int_F \sigma_j(t_j, \theta)[a_j] \times \pi_i^*(t_i)[dt_j, \theta]$.

Definition 3 *The correspondence of best replies for all types given a subset of actions for all types is denoted $BR : \left((2^{A_i})^{T_i^*} \right)_{i \in I} \rightarrow \left((2^{A_i})^{T_i^*} \right)_{i \in I}$ and is defined as follows. First, given a specification of a subset of actions for each possible type of opponent, denoted by $E_j = \left((E_{t_j})_{t_j \in T_j^*} \right)_{j \neq i}$, with $E_{t_j} \subset A_j$ for all t_j and $j \neq i$, we define the ε best replies for t_i in game G as*

$$BR_i(t_i, E_j) = \left\{ a_i \in A_i \left| \begin{array}{l} \exists \nu \in \Delta(T_j^* \times \Theta \times A_j) \text{ such that} \\ (i) \nu[\{(t_j, \theta, a_j) : a_j \in E_j\}] = 1 \\ (ii) \text{marg}_{T_j^* \times \Theta} \nu = \pi_i^*(t_i) \\ (iii) \int_{(t_j, \theta, a_j)} \begin{bmatrix} g_i(a_i, a_j, \theta) \\ -g_i(a'_i, a_j, \theta) \end{bmatrix} d\nu \geq -\varepsilon \text{ for all } a'_i \in A_i \end{array} \right. \right\}$$

Remark 1 *In cases where E_j is not measurable, we interpret $\nu[\{(t_j, \theta, a_j) : a_j \in E_j\}] = 1$ as saying that there is a measurable subset $E' \subseteq E_j$ such that $\nu[\Theta \times E'] = 1$.*

The solution concept and closely related notions with which we work in this paper are given below.

Definition 4 *1. The interim rationalizable set, $R = \left((R_i(t_i))_{t_i \in T_i^*} \right)_{i \in I} \subset (A^{T_i^*})_{i \in I}$, is the largest fixed point of BR.*

2. $R_0 \equiv \Pi_i(A_i^{T_i^})$, $R_k \equiv BR(R_{k-1})$, and $R_\infty \equiv \bigcap_{k=1}^\infty R_k$.*

3. Let $S = (S_1, S_2)$, where each $S_i : T^ \rightarrow 2^{A_i} / \emptyset$; S is a best-reply set if for each t_i and $a_i \in S_i(t_i)$, there exists a measurable $\nu \in \Delta(T_j^* \times \Theta \times A_j)$ such that*

$$\begin{array}{l} (i) \nu[\{(t_j, \theta, a_j) : a_j \in S_j\}] = 1 \\ (ii) \text{marg}_{T_j^* \times \Theta} \nu = \pi_i^*(t_i) \\ (iii) \int_{(t_j, \theta, a_j)} \begin{bmatrix} g_i(a_i, a_j, \theta) \\ -g_i(a'_i, a_j, \theta) \end{bmatrix} d\nu \geq -\varepsilon \text{ for all } a'_i \in A_i \end{array}$$

Arguments in Dekel, Fudenberg and Morris (2005) establish that these sets are well-defined, and the relationships among them. These are restated in the following results that are taken from that work.

- Result 1**
1. $R_{i,k}(t_i) = R_{i,k}(t'_i)$ if $\delta_k(t_i) = \delta_k(t'_i)$. That is, types with the same k^{th} -order beliefs have the same k^{th} -order rationalizable sets.
 2. If S_c^T for all c in some index set C are best-reply sets then $\cup_c S_c^T$ is a best-reply set.
 3. The union of all best-reply sets is a best reply set. It is also the largest fixed point of BR.
 4. R equals R_∞ .
 5. $R_{i,k}$ and $R_{i,\infty}$ are measurable functions from $T_i \rightarrow 2^{A_i}/\emptyset$, and for each action a_i and each k the sets $\{t : a_i \in R_{i,k}(t)\}$ and $\{t : a_i \in R_{i,\infty}(t)\}$ are closed.

In defining our strategic topology, we will exploit the following closure properties of ε -rationalizable sets as a function of ε :

Lemma 1 For each $k = 0, 1, \dots$, if $\varepsilon^n \downarrow \varepsilon$ and $a_i \in R_i^k(t_i, G, \varepsilon^n)$ for all n , then $a_i \in R_i^k(t_i, G, \varepsilon)$.

Proof. We will prove this by induction. It is vacuously true for $k = 0$

Suppose that it holds true up to $k - 1$. Let

$$\Psi_i^k(t_i, \delta) = \left\{ \psi \in \Delta(A_j \times \Theta) : \begin{array}{l} \psi(a_j, \theta) = \int_{T_j^*} \nu [dt_j, \theta, a_j] \\ \text{for some } \nu \in \Delta(T_j^* \times \Theta \times A_j) \text{ such that} \\ \nu[\{(t_j, \theta, a_j) : a_j \in R_j^{k-1}(t_j, G, \varepsilon)\}] = 1 \\ \text{and } \text{marg}_{T^* \times \Theta} \nu = \pi_i^*(t_i) \end{array} \right\}.$$

The sequence $\Psi_i^k(t_i, \varepsilon^n)$ is decreasing (under set inclusion) and converges (by σ -additivity) to $\Psi_i^k(t_i, \varepsilon)$; moreover, in Dekel, Fudenberg and Morris (2005) we show that each $\Psi_i^k(t_i, \varepsilon^n)$ is compact. Let

$$\Lambda_i^k(t_i, a_i, \delta) = \left\{ \psi \in \Delta(A_j \times \Theta) : \sum_{a_j, \theta} \psi(a_j, \theta) \begin{bmatrix} g_i(a_i, a_j, \theta) \\ -g_i(a'_i, a_j, \theta) \end{bmatrix} \geq -\delta \text{ for all } a'_i \in A_i \right\}.$$

The sequence $\Lambda_i^k(t_i, a_i, \varepsilon^n)$ is decreasing (under set inclusion) and converges to $\Lambda_i^k(t_i, a_i, \varepsilon)$.
Now $a_i \in R_i^k(t_i, G, \varepsilon^n)$ for all n

$$\begin{aligned} &\Rightarrow \Psi_i^k(t_i, \varepsilon^n) \cap \Lambda_i^k(t_i, a_i, \varepsilon^n) \neq \emptyset \text{ for all } n \\ &\Rightarrow \Psi_i^k(t_i, \varepsilon) \cap \Lambda_i^k(t_i, a_i, \varepsilon) \neq \emptyset \\ &\Rightarrow a_i \in R_i^k(t_i, G, \varepsilon), \end{aligned}$$

where the second implication follows from the finite intersection property of compact sets. ■

Proposition 1 *If $\varepsilon^n \downarrow \varepsilon$ and $a_i \in R_i(t_i, G, \varepsilon^n)$ for all n , then $a_i \in R_i(t_i, G, \varepsilon)$.*

Proof.

$$\begin{aligned} a_i &\in R_i(t_i, G, \varepsilon^n) \text{ for all } n \\ &\Rightarrow a_i \in R_i^k(t_i, G, \varepsilon^n) \text{ for all } n \text{ and } k \\ &\Rightarrow a_i \in R_i^k(t_i, G, \varepsilon) \text{ for all } k, \text{ by the above lemma} \\ &\Rightarrow a_i \in R_i(t_i, G, \varepsilon). \end{aligned}$$

■

Corollary 1 *For any t_i , a_i and G*

$$\min \{ \varepsilon : a_i \in R_i(t_i, G, \varepsilon) \}$$

exists.

5 The Strategic Topology

Our goal is to find the coarsest topology on types so that the ε -best-response correspondence and ε -rationalizable sets have continuity properties such as those satisfied by ε -Nash equilibrium and ε -rationalizability in complete information games with respect to the payoffs. For any fixed game and action, we define the distance between a pair of types as the difference between the smallest ε that would make that action ε -rationalizable in that game.

Thus

$$\begin{aligned} h_i(t_i|a_i, G) &= \min \{ \varepsilon : a_i \in R_i(t_i, G, \varepsilon) \} \\ d(t_i, t'_i|a_i, G) &= |h_i(t_i|a_i, G) - h_i(t'_i|a_i, G)| \end{aligned}$$

We will write \mathcal{G}^m for the collection of games where each player has at most m actions.

Definition 5 (Upper Strategic Convergence) $((t_i^n)_{n=1}^\infty, t_i)$ satisfy the upper strategic convergence property (written $t_i^n \rightarrow_U t_i$) if $\forall m, \exists \varepsilon^n \rightarrow 0$ s.t. $h_i(t_i|a_i, G) < h_i(t_i^n|a_i, G) + \varepsilon^n$ for all $n, a_i, G \in \mathcal{G}^m$.

This implies that if $\varepsilon^n \rightarrow 0$ and $a_i \in R_i(t_i^n, G, \varepsilon^n)$ for each n , then $a_i \in R_i(t_i, G, 0)$.¹²

We do not require convergence uniformly over all games, since an upper bound on the number of actions m is fixed before the approximating sequence ε^n is chosen. Requiring uniformity over all games would considerably strengthen the topology, as briefly discussed in Section 7.2. On the other hand, it will be clear from our arguments that our results would not be changed if we dropped the "partial" uniformity in the definition and simply required pointwise convergence, i.e., for all $a_i, G \in \mathcal{G}$, there exists $\varepsilon^n \rightarrow 0$ s.t. $h_i(t_i|a_i, G) < h_i(t_i^n|a_i, G) + \varepsilon^n$. But it is more convenient to state the upper strategic convergence property in the form of our definition.

Definition 6 (Lower Strategic Convergence) $((t_i^n)_{n=1}^\infty, t_i)$ satisfy the lower strategic convergence property (written $t_i^n \rightarrow_L t_i$) if $\forall m, \exists \varepsilon^n \rightarrow 0$ s.t. $h_i(t_i^n|a_i, G) < h_i(t_i|a_i, G) + \varepsilon^n$ for all $n, a_i, G \in \mathcal{G}^m$.

This implies that if $a_i \in R_i(t_i, G, 0)$, then there exists $\varepsilon^n \rightarrow 0$ such that $a_i \in R_i(t_i^n, G, \varepsilon^n)$ for each n . This is simply the mirror image of the upper strategic convergence property and requires that $h_i(t_i^n|a_i, G)$ does not drop in the limit for all values of $h_i(t_i|a_i, G)$ (not just for $h_i(t_i|a_i, G) = 0$).

Now consider the following notion of distance between types:

$$d(t_i, t_i') = \sum_m \beta^m \sup_{a_i, G \in \mathcal{G}^m} d(t_i, t_i'|a_i, G) \quad (1)$$

where $0 < \beta < 1$.

Lemma 2 *The distance d is a metric.*

¹²It is possibly stronger than this property since upper convergence requires that $h_i(t_i^n|a_i, G)$ does not jump in the limit for all values of $h_i(t_i|a_i, G)$ (not just for $h_i(t_i|a_i, G) = 0$). We have not verified if it is strictly stronger.

Proof. First note d is symmetric by definition. To see that d satisfies the triangle inequality, note that for each action a_i and game G ,

$$\begin{aligned} d(t_i, t_i'' | a_i, G) &= |h_i(t_i | a_i, G) - h_i(t_i'' | a_i, G)| \\ &\leq |h_i(t_i | a_i, G) - h_i(t_i' | a_i, G)| + |h_i(t_i' | a_i, G) - h_i(t_i'' | a_i, G)| \\ &= d(t_i, t_i' | a_i, G) + d(t_i', t_i'' | a_i, G) \end{aligned}$$

hence

$$\begin{aligned} d(t_i, t_i'') &= \sum_m \beta^m \sup_{a_i, G \in \mathcal{G}^m} d(t_i, t_i'' | a_i, G) \\ &\leq \sum_m \beta^m \sup_{a_i, G \in \mathcal{G}^m} (d(t_i, t_i' | a_i, G) + d(t_i', t_i'' | a_i, G)) \\ &\leq \sum_m \beta^m \left(\sup_{a_i, G \in \mathcal{G}^m} d(t_i, t_i' | a_i, G) + \sup_{a_i, G \in \mathcal{G}^m} d(t_i', t_i'' | a_i, G) \right) \\ &= \sum_m \beta^m \sup_{a_i, G \in \mathcal{G}^m} d(t_i, t_i' | a_i, G) + \sum_m \beta^m \sup_{a_i, G \in \mathcal{G}^m} d(t_i', t_i'' | a_i, G) \\ &= d(t_i, t_i') + d(t_i', t_i''). \end{aligned}$$

Theorem 1 below implies that $d(t_i, t_i') = 0 \Rightarrow t_i = t_i'$. ■

Lemma 3 $d(t_i^n, t_i) \rightarrow 0$ if and only if $t_i^n \rightarrow_U t_i$ and $t_i^n \rightarrow_L t_i$.

Proof. Suppose $d(t_i^n, t_i) \rightarrow 0$. Fix m and let

$$\varepsilon^n = \beta^{-m} d(t_i^n, t_i).$$

Now for any a_i and $G \in \mathcal{G}^m$,

$$\beta^m |h_i(t_i^n | a_i, G) - h_i(t_i | a_i, G)| \leq \sum_m \beta^m \sup_{a_i', G' \in \mathcal{G}^m} d(t_i^n, t_i | a_i', G') = d(t_i^n, t_i);$$

so

$$|h_i(t_i^n | a_i, G) - h_i(t_i | a_i, G)| \leq \beta^{-m} d(t_i^n, t_i) = \varepsilon^n;$$

thus $t_i^n \rightarrow_U t_i$ and $t_i^n \rightarrow_L t_i$.

Conversely, suppose that $t_i^n \rightarrow_U t_i$ and $t_i^n \rightarrow_L t_i$. Then $\forall m, \exists \bar{\varepsilon}^n(m) \rightarrow 0$ and $\underline{\varepsilon}^n(m) \rightarrow 0$ s.t. for all $a_i, G \in \mathcal{G}^m$,

$$\begin{aligned} h_i(t_i | a_i, G) &< h_i(t_i^n | a_i, G) + \bar{\varepsilon}^n(m) \\ \text{and } h_i(t_i^n | a_i, G) &< h_i(t_i | a_i, G) + \underline{\varepsilon}^n(m). \end{aligned}$$

Thus

$$\begin{aligned}
d(t_i^n, t_i) &= \sum_m \beta^m \sup_{a_i, G \in \mathcal{G}^m} d(t_i^n, t_i | a_i, G) \\
&\leq \sum_m \beta^m \max(\bar{\varepsilon}^n(m), \underline{\varepsilon}^n(m)) \\
&\rightarrow 0 \text{ as } n \rightarrow \infty.
\end{aligned}$$

■

Definition 7 *The strategic topology is the topology generated by metric d .*

The strategic topology is thus a metric topology where sequences converge if and only if they satisfy upper and lower strategic convergence. In general, the closed sets in a topology are not determined by the convergent sequences, and extending the convergence definitions to a topology can introduce more convergent sequences. However, since metric spaces are first countable, convergence and continuity can be assessed by looking at sequences (Munkres (1975), p.190), and so this must also be the coarsest metric topology with the desired continuity properties. We now illustrate the strategic topology with the example discussed earlier.

5.1 The E-Mail Example Revisited

We can illustrate the definitions in this section with the e-mail example introduced informally earlier. We show that we have convergence of types in the product topology, $t_{1k} \rightarrow^* t_{1\infty}$, corresponding to the upper-hemicontinuity noted in section 2, while $d(t_{1k}, t_{1\infty}) \not\rightarrow 0$, corresponding to the failure of lower hemi-continuity.

Writing \widehat{G} for the payoffs in the example, we observe that beliefs can be defined as follows:

$$\begin{aligned}
\pi^*(t_{11})[(t_2, \theta)] &= \begin{cases} 1, & \text{if } (t_2, \theta) = (t_{21}, 0) \\ 0, & \text{otherwise} \end{cases} \\
\pi^*(t_{1m})[(t_2, \theta)] &= \begin{cases} \frac{1}{2-\alpha}, & \text{if } (t_2, \theta) = (t_{2,m-1}, 1) \\ \frac{1-\alpha}{2-\alpha}, & \text{if } (t_2, \theta) = (t_{2m}, 1) \\ 0, & \text{otherwise} \end{cases}, \text{ for all } m = 2, 3, \dots \\
\pi^*(t_{1\infty})[(t_2, \theta)] &= \begin{cases} 1, & \text{if } (t_2, \theta) = (t_{2\infty}, 1) \\ 0, & \text{otherwise} \end{cases}
\end{aligned}$$

$$\begin{aligned}
\pi^*(t_{21})[(t_1, \theta)] &= \begin{cases} \frac{1}{2-\alpha}, & \text{if } (t_1, \theta) = (t_{11}, 0) \\ \frac{1-\alpha}{2-\alpha}, & \text{if } (t_2, \theta) = (t_{12}, 1) \\ 0, & \text{otherwise} \end{cases} \\
\pi^*(t_{2m})[(t_1, \theta)] &= \begin{cases} \frac{1}{2-\alpha}, & \text{if } (t_1, \theta) = (t_{1m}, 1) \\ \frac{1-\alpha}{2-\alpha}, & \text{if } (t_2, \theta) = (t_{1,m+1}, 1) \quad , \text{ for all } m = 2, 3, \dots \\ 0, & \text{otherwise} \end{cases} \\
\pi^*(t_{2\infty})[(t_1, \theta)] &= \begin{cases} 1, & \text{if } (t_1, \theta) = (t_{1\infty}, 1) \\ 0, & \text{otherwise} \end{cases}
\end{aligned}$$

Thus we have product convergence, $t_{1k} \rightarrow t_{1\infty}$.

Now for any $\varepsilon < \frac{1+\alpha}{2-\alpha}$,

$$\begin{aligned}
R_1^0(t_1, \widehat{G}, \varepsilon) &= \{N, I\} \text{ for all } t_1 \\
R_2^0(t_2, \widehat{G}, \varepsilon) &= \{N, I\} \text{ for all } t_2
\end{aligned}$$

$$\begin{aligned}
R_1^1(t_1, \widehat{G}, \varepsilon) &= \begin{cases} \{N\}, & \text{if } t_1 = t_{11} \\ \{N, I\}, & \text{if } t_1 \in \{t_{12}, t_{13}, \dots\} \cup \{t_{1\infty}\} \end{cases} \\
R_2^1(t_2, \widehat{G}, \varepsilon) &= \{N, I\}
\end{aligned}$$

$$\begin{aligned}
R_1^2(t_1, \widehat{G}, \varepsilon) &= \begin{cases} \{N\}, & \text{if } t_1 = t_{11} \\ \{N, I\}, & \text{if } t_1 \in \{t_{12}, t_{13}, \dots\} \cup \{t_{1\infty}\} \end{cases} \\
R_2^2(t_2, \widehat{G}, \varepsilon) &= \begin{cases} \{N\}, & \text{if } t_2 = t_{21} \\ \{N, I\}, & \text{if } t_2 \in \{t_{22}, t_{23}, \dots\} \cup \{t_{2\infty}\} \end{cases}
\end{aligned}$$

$$\begin{aligned}
R_1^3(t_1, \widehat{G}, \varepsilon) &= \begin{cases} \{N\}, & \text{if } t_1 \in \{t_{11}, t_{12}\} \\ \{N, I\}, & \text{if } t_1 \in \{t_{13}, t_{14}, \dots\} \cup \{t_{1\infty}\} \end{cases} \\
R_2^3(t_2, \widehat{G}, \varepsilon) &= \begin{cases} \{N\}, & \text{if } t_2 = t_{21} \\ \{N, I\}, & \text{if } t_2 \in \{t_{22}, t_{23}, \dots\} \cup \{t_{2\infty}\} \end{cases}
\end{aligned}$$

$$\begin{aligned}
R_1^{2m}(t_1, \widehat{G}, \varepsilon) &= \begin{cases} \{N\}, & \text{if } t_1 \in \{t_{11}, \dots, t_{1m}\} \\ \{N, I\}, & \text{if } t_1 \in \{t_{1,m+1}, t_{1,m+2}, \dots\} \cup \{t_{1\infty}\} \end{cases} \\
R_2^{2m}(t_2, \widehat{G}, \varepsilon) &= \begin{cases} \{N\}, & \text{if } t_2 \in \{t_{21}, \dots, t_{2m}\} \\ \{N, I\}, & \text{if } t_2 \in \{t_{2,m+1}, t_{2,m+2}, \dots\} \cup \{t_{2\infty}\} \end{cases}
\end{aligned}$$

for $m = 2, 3, \dots$

$$R_1^{2m+1}(t_1, \widehat{G}, \varepsilon) = \begin{cases} \{N\}, & \text{if } t_1 \in \{t_{11}, \dots, t_{1,m+1}\} \\ \{N, I\}, & \text{if } t_1 \in \{t_{1,m+2}, t_{1,m+3}, \dots\} \cup \{t_{1\infty}\} \end{cases}$$

$$R_2^{2m+1}(t_2, \widehat{G}, \varepsilon) = \begin{cases} \{N\}, & \text{if } t_2 \in \{t_{21}, \dots, t_{2m}\} \\ \{N, I\}, & \text{if } t_2 \in \{t_{2,m+1}, t_{2,m+2}, \dots\} \cup \{t_{2\infty}\} \end{cases}$$

for $m = 2, 3, \dots$; so

$$R_1(t_1, \widehat{G}, \varepsilon) = \begin{cases} \{N\}, & \text{if } t_1 \in \{t_{11}, t_{12}, \dots\} \\ \{N, I\}, & \text{if } t_1 = t_{1\infty} \end{cases}$$

$$R_2(t_2, \widehat{G}, \varepsilon) = \begin{cases} \{N\}, & \text{if } t_2 \in \{t_{21}, t_{22}, \dots\} \\ \{N, I\}, & \text{if } t_2 = t_{2\infty} \end{cases}$$

Now observe that

$$h_1(t_{1k} | I, \widehat{G}) = \begin{cases} 2, & \text{if } k = 1 \\ \frac{1+\alpha}{2-\alpha}, & \text{if } k = 2, 3, \dots \\ 0, & \text{if } k = \infty \end{cases}$$

(while $h_1(t_{1k} | N, \widehat{G}) = 0$ for all k). Thus $d(t_{1k}, t_{1\infty}) \geq \frac{1+\alpha}{2-\alpha}$ for all $k = 2, 3, \dots$ and we do not have $d(t_{1k}, t_{1\infty}) \rightarrow 0$.

6 Results

6.1 The relationships among the notions of convergence

We first demonstrate that both lower strategic convergence and upper strategic convergence imply product convergence.

Theorem 1 *Upper strategic convergence implies product convergence. Lower strategic convergence implies product convergence.*

These results follow from a pair of lemmas. The product topology is generated by the metric

$$\widetilde{d}(t_i, t'_i) = \sum_k \beta^k \widetilde{d}^k(t_i, t'_i)$$

where $0 < \beta < 1$ and \tilde{d}^k is a metric on the k th level beliefs that generates the topology of weak convergence. One such metric is the Prokhorov metric, which is defined as follows. For any metric space X , let \mathcal{F} be the Borel sets, and for $A \in \mathcal{F}$ set $A^\gamma = \{x \in X \mid \inf_{y \in A} |x - y| \leq \gamma\}$. Then the Prokhorov distance between measures δ and δ' is $\pi(\delta, \delta') = \inf\{\gamma \mid \delta(A) \leq \delta'(A^\gamma) + \gamma \text{ for all } A \in \mathcal{F}\}$, and $\tilde{d}^k(t_i, t'_i) = \pi(\delta_k(t_i), \delta_k(t'_i))$.

Lemma 4 *For all k and $c > 0$, there exist $\varepsilon > 0$ and m s.t. if $\tilde{d}^k(t_i, t'_i) > c$, $\exists a_i, G \in \mathcal{G}^m$ s.t. $h_i(t'_i | a_i, G) + \varepsilon < h_i(t_i | a_i, G)$.*

Proof. To prove this we will construct a variant of a "report your beliefs" game, and show that any two types whose k^{th} order beliefs differ by δ will lose a non-negligible amount by taking the (unique) rationalizable action of the other type.

To define the finite games we will use for the proof, it is useful to first think of a very large infinite action game where the action space is the type space T^* . Thus the first component of player i 's action is a probability distribution over Θ : $a_i^1 \in \Delta(\Theta)$. The second component of the action is an element of $\Delta(\Theta \times \Delta(\Theta))$, and so on. The idea of the proof is to start with a proper scoring rule for this infinite game (so that each player has a unique rationalizable action, which is to truthfully report his type), and use it to define a finite game where the rationalizable actions are "close to truth telling."

To construct the finite game, we have agents report only the first k levels of beliefs, and impose a finite grid on the reports at each level. Specifically, for any fixed integer z^1 let A^1 be the set of probability distributions a^1 on Θ such that for all $\theta \in \Theta$, $a^1(\theta) = j/z^1$ for some integer j , $1 \leq j \leq z^1$. Thus $A^1 = \{a \in \mathbf{R}^{|\Theta|} : a_\theta = j/z^1 \text{ for some integer } j, 1 \leq j \leq z^1, \sum_\theta a_\theta = 1\}$; it is a discretization of the set $\Delta(\Theta)$ with grid points that are evenly spaced in the Euclidean metric.

Let $D^1 = \Theta \times A^1$. Note that this is a finite set. Next pick an integer z^2 and let A^2 be the set of probability distributions on D^1 such that $a^2(d) = j/z^2$ for some integer j , $1 \leq j \leq z^2$. Continuing in this way we can define a sequence of finite action sets A^j , where every element of each A^j is a probability distribution with finite support. The overall action chosen is a vector in $A^1 \times A^2 \times \dots \times A^k$.

We call the a^m the " m^{th} -order action." Let the payoff function be

$$g_i(a_1, a_2, \theta) = 2a_i^1(\theta) - \sum_{\theta'} (a_i^1(\theta'))^2 + \sum_{m=2}^k \left[2a_i^m(a_j^1, \dots, a_j^{m-1}, \theta) - \sum_{\tilde{a}_j^1, \dots, \tilde{a}_j^{m-1}, \tilde{\theta}} \left(a_i^m(\tilde{a}_j^1, \dots, \tilde{a}_j^{m-1}, \tilde{\theta}) \right)^2 \right].$$

Note that the objective functions are strictly concave, and that the payoff to the m^{th} -order action depends only on the state θ and on actions of the other player up to the $(m - 1)^{\text{th}}$ level (so the payoff to a_i^1 does not depend on player j 's action at all). This will allow us to determine the rationalizable sets recursively, starting from the first-order actions and working up.

The rationalizable first-order action(s) for type t_i with first-order beliefs δ_1 is the point or points $a_1^*(\delta_1) \in A^1$ that is closest to δ_1 . Picking any other point involves a loss, so for each grid size z^1 there is an $\varepsilon > 0$ such that no other actions are ε -rationalizable for t . Moreover, any type whose first-order beliefs are sufficiently different from $\delta_1(t)$ will want to pick a different action. Thus for all $c > 0$, if $\tilde{d}(\delta_1(t_i), \delta_1(t'_i)) > c$, there are $\varepsilon_1 > 0$ and z^1 such that

$$h_i(t'_i | a_1^*(t'_i), G) + \varepsilon_1 < h_i(t_i | \delta_1^*(t'_i), G).$$

This proves the claim for the case $k = 1$.

Now let $\delta_2(t_i) \in \Delta(\Theta \times \Delta(\Theta))$ be the second-order belief of t_i . For any fixed first-level grid z^1 , we know from the first step that there is an $\varepsilon_1 > 0$ such that for any δ_1 , the only ε_1 -rationalizable first-order actions are the point or points a_1^* in the grid that are closest to δ_1 . Suppose that player i believes player j is playing a first-order action that is ε_1 -rationalizable. Then player i 's beliefs about the finite set $D^1 = \Theta \times A^1$ correspond to the finite-dimensional measure δ_2^* , where the probability of measurable $X \subseteq \Theta \times A_1$ is $\delta_2^*(X) = \delta_2(\{(\delta_1, \theta) : (\theta, a_1^*(\delta_1)) \in X\})$. That is, for each δ_1 that i thinks j could have, i expects that j will play an element of the corresponding $a_1^*(\delta_1)$. Because A^2 is a discretization of $\Delta(D^1)$, player i may not be able to choose $a_2 = \delta_2^*$. However, because of the concavity of the objective function, the constrained second-order best reply of i with beliefs δ_2 is the point $a_2^* \in A_2$ that is closest to δ_2^* in the Euclidean metric, and choosing any other action incurs a non-zero loss. Moreover, a_2^* is at (Euclidean) distance from δ_2^* that is bounded by the distance between grid points, so there is a bound on the distance that goes to zero as z^2 goes to infinity, uniformly over all δ_2^* .

Next we claim that if there is a $c > 0$ such that $d^2(t_i, t'_i) > c$, then $\delta_2^*(t_i) \neq \delta_2^*(t'_i)$ for all sufficiently fine grids A^1 in $\Delta(\Theta^1)$. To see this, note from the definition of the Prokhorov metric, if $\tilde{d}^2(t_i, t'_i) > c$ there is a Borel set A in $\Theta \times \Delta(\Theta)$ such that $\delta_2(t_i)(A) > \delta_2(t'_i)(A^c) + c$. Because the first-order actions a_1^* converge uniformly to δ_1^* as z^2 goes to infinity, $(\theta, a_1^*(\delta_1)) \in A^c$ for every $(\theta, \delta_1) \in A$, so for all γ such that $c/2 > \gamma > 0$ there is a \bar{z}_2 such that $\delta_2^*(t_i)(A^c) \geq$

$\delta_2^*(t_i)(A) > \delta_2(t_i)(A) - \gamma > \delta_2(t'_i)(A^c) - \gamma + c \geq \delta_2^*(t'_i)(A^c) - 2\gamma + c > \delta_2^*(t'_i)(A^c)$, where the first inequality follows from set inclusion, the second and fourth from the uniform convergence of the a_1^* , and the third from $\tilde{d}^2(t_i, t'_i) > c$.

From the previous step, this implies that when $d^2(t_i, t'_i) > c$ there is a \bar{z}^2 and $\varepsilon_2 > 0$ such that

$$h_i(t'_i | a_2^*(t'_i), G) + \varepsilon_2 < h_i(t_i | \delta_2^*(t'_i), G)$$

for all $z_2 > \bar{z}_2$.

We can continue in this way to prove the result for any k . ■

Lemma 5 *Suppose that t_i is not the limit of the sequence t_i^n in the product topology, then (t_i^n, t_i) satisfies neither the lower convergence property nor the upper convergence property.*

Proof. Failure of product convergence implies that there exists k such that $\tilde{d}^k(t_i^n, t_i)$ does not converge to zero, so there exists $\delta > 0$ such that (for some subsequence)

$$\tilde{d}^k(t_i^n, t_i) > \delta$$

for all n .

By lemma 4, there exists ε and m such that, for all n ,

- (a) $\exists a_i, G \in \mathcal{G}^m$ s.t. $h_i(t_i | a_i, G) + \varepsilon < h_i(t_i^n | a_i, G)$ and
- (b) $\exists a_i, G \in \mathcal{G}^m$ s.t. $h_i(t_i^n | a_i, G) + \varepsilon < h_i(t_i | a_i, G)$.

Now suppose that the lower convergence property holds. Then there exists $\eta^n \rightarrow 0$ such that

$$h_i(t_i^n | a_i, G) < h_i(t_i | a_i, G) + \eta^n$$

for all $a_i, G \in \mathcal{G}^m$. This combined with (a) gives a contradiction.

Similarly, upper convergence implies that there exists $\eta^n \rightarrow 0$ such that

$$h_i(t_i | a_i, G) < h_i(t_i^n | a_i, G) + \eta^n$$

for all $a_i, G \in \mathcal{G}^m$. This give a contradiction when combined with (b). ■

Lemma 5 immediately implies theorem 1.

Theorem 2 *Product convergence implies upper strategic convergence.*

Proof. Suppose that t_i^n product-converges to t_i . If upper strategic convergence fails there is an $m, a_i, G \in \mathcal{G}^m$ s.t. for all $\varepsilon^n \rightarrow 0$ and N , there is $n' > N$ such that

$$h_i(t_i|a_i, G) > h_i(t_i^{n'}|a_i, G) + \varepsilon^{n'}.$$

We may relabel so that t_i^n is the subsequence where this inequality holds. Pick δ so that

$$h_i(t_i|a_i, G) > h_i(t_i^n|a_i, G) + \delta$$

for all n . Since, for each n and t_i^n , $a_i \in R_i(t_i^n, G, h_i(t_i|a_i, G) - \delta)$, there exists $\nu^n \in \Delta(T_j^* \times \Theta \times A_j)$ s.t.

- (i) $\nu^n [\{(t_j, \theta, a_j) : a_j \in R_j(t_j, G, h_i(t_i|a_i, G) - \delta) \text{ for all } j \neq i\}] = 1$
- (ii) $\text{marg}_{T^* \times \Theta} \nu^n = \pi^*(t_i^n)$
- (iii) $\int_{(t_j, \theta, a_j)} \begin{bmatrix} g_i(a_i, a_j, \theta) \\ -g_i(a'_i, a_j, \theta) \end{bmatrix} d\nu^n \geq -h_i(t_i|a_i, G) + \delta$ for all $a'_i \in A_i$

Since under the product topology, T^* is a compact metric space, and since A_j and Θ are finite, so is $T^* \times \Theta \times A_j$. Thus $\Delta(T^* \times \Theta \times A_j)$ is compact in the weak topology, so the sequence ν^n has a limit point, ν .

Now since (i), (ii) and (iii) holds for every n and $\nu = \lim_n \nu^n$, we have

- (i*) $\nu [\{(t_j, \theta, a_j) : a_j \in R_j(t_j, G, h_i(t_i|a_i, G) - \delta)\}] = 1$
- (ii*) $\text{marg}_{T^* \times \Theta} \nu = \pi^*(t_i)$
- (iii*) $\int_{(t_j, \theta, a_j)} \begin{bmatrix} g_i(a_i, a_j, \theta) \\ -g_i(a'_i, a_j, \theta) \end{bmatrix} d\nu \geq -h_i(t_i|a_i, G) + \delta$ for all $a'_i \in A_i$

where (i*) follows from the fact that $\{(t_j, \theta, a_j) : a_j \in R_j(t_j, G, h_i(t_i|a_i, G) - \delta)\}$ is closed. This implies $a_i \in R_i(t_i, G, h_i(t_i|a_i, G) - \delta)$, a contradiction. ■

Corollary 2 *Lower strategic convergence implies convergence in the strategic topology.*

Proof. We have $t_i^n \rightarrow_L t_i \Rightarrow t_i^n \rightarrow^* t_i$ (by theorem 1); $t_i^n \rightarrow^* t_i \Rightarrow t_i^n \rightarrow_U t_i$ (by theorem 2); and $t_i^n \rightarrow_L t_i$ and $t_i^n \rightarrow_U t_i \Rightarrow d(t_i^n, t_i) \rightarrow 0$ (by lemma 3). ■

6.2 Finite types are dense in the strategic topology

Theorem 3 *Finite types are dense under d .*

Given corollary 2, the theorem follows from lemma 6 below, which shows that, for any type in the universal type space, it is possible to construct a sequence of finite types that lower converge to it. The proof of lemma 6 is long, and broken into many steps. In outline, we first find a finite grid of games that approximate all games with m actions. We then map types onto their "epsilons" for action/game pairs from this finite set, impose a grid on these epsilons, and define a belief hierarchy for a finite type by arbitrary taking the belief hierarchy of one of the types that is mapped to it. This gives us a finite type space. We show that this map "preserve the epsilons": $a_i \in R_i(t_i, g, \varepsilon) \Rightarrow a_i \in R_i(f_i(t_i), g, \varepsilon)$. Finally we show that for any type in the universal space there is a sequence of these finite types that "lower converge" to it. Our proof thus follows Monderer and Samet (1996) in constructing a mapping from types in one type space to types in another type space that preserves approximate best response properties. Their construction worked for equilibrium, while our construction works for rationalizability, and thus the approximation has to work for many conjectures over opponents' play simultaneously. We assume neither a common prior nor a countable number of types, and we develop a topology on types based on the play of the given types as opposed to a topology on priors or information systems.¹³

A distinctive feature of our work is that we identify types in our constructed type space with sets of ε -rationalizable actions for a finite set of ε and a finite set of games. The recent paper of Ely and Peski (2005) similarly identifies types with sets of rationalizable actions, although for their different purpose (constructing a universal type space for the independent interim rationalizability solution concept), no approximation is required.

Lemma 6 *For any t_i , there exists a sequence of finite types t_i^n such that $[(t_i^n)_{n=1}^\infty, t_i]$ satisfy lower strategic convergence.*

For the first part of the proof, we will restrict attention to m action games with $A_1 = A_2 = \{1, 2, \dots, m\}$. Having fixed the action sets, a game is parameterized by the payoff

¹³It is not clear how one could develop a topology based on the equilibrium distribution of play in a setting without a common prior. At a technical level, the fact that the strategic topology is not uniform over the number of actions requires extra steps to deal with the cardinality of the number of actions. (See Section 7.2 for a brief discussion of uniformity.)

function $g : A_1 \times A_2 \times \Theta \rightarrow [-M, M]^2$. For a fixed m , we will write g for the game $G = (\{1, 2, \dots, m\}, \{1, 2, \dots, m\}, g)$. Let

$$D(g, g') = \sup_{i, a, \theta} |g_i(a, \theta) - g'_i(a, \theta)|.$$

Some preliminary lemmas are proved in the Appendix. There is a finite grid of games with m actions such all games are within ε of a game in the grid:

Lemma 7 *For any integer m and $\varepsilon > 0$, there exists a finite collection of m action games $\bar{\mathcal{G}}$ such that, for every $g \in \mathcal{G}^m$, there exists $g' \in \bar{\mathcal{G}}$ such that $D(g, g') \leq \varepsilon$.*

The rationalizability of an action ($h_i(t_i|a_i, g)$) is close to the rationalizability of that action in nearby games:

Lemma 8 *For all i, t_i, a_i, g and g' ,*

$$h_i(t_i|a_i, g) \leq h_i(t_i|a_i, g') + 2D(g, g')$$

for all i, a_i .

So the rationalizability of an action is close to the rationalizability of that action in a nearby game in the grid:

Lemma 9 *For any integer m and $\varepsilon > 0$, there exists a finite collection of m action games $\bar{\mathcal{G}}$ such that, for every $g \in \mathcal{G}^m$, there exists $g' \in \bar{\mathcal{G}}$ such that*

$$|h_i(t_i|a_i, g) - h_i(t_i|a_i, g')| \leq \varepsilon.$$

for all i, t_i, a_i, g and g' .

The main step of the proof of lemma 6 is then the construction of a finite type space that replicates the ε -rationalizability properties all games in the grid of m action games if ε is in some finite grid.

Lemma 10 *Fix any finite collection of m action games $\bar{\mathcal{G}}$ and $\delta > 0$. There exists a finite type space $(T_i, \hat{\pi}_i)_{i=1,2}$ and functions $(f_i)_{i=1,2}$, each $f_i : T^* \rightarrow T_i$, such that $R_i(t_i, g, \varepsilon) \subseteq R_i(f_i(t_i), g, \varepsilon)$ for all $t_i \in T^*$ and $\varepsilon \in \{0, \delta, 2\delta, \dots\}$.*

Proof. Write $\langle x \rangle^\delta$ for the smallest number in the set $\Delta = \{0, \delta, 2\delta, \dots\}$ greater than x . Let

$$f_i : T^* \rightarrow \left\{ \widehat{t}_i : \overline{\mathcal{G}} \times A \rightarrow \left\{ 0, \delta, \dots, \langle 2M \rangle^\delta \right\} \right\}$$

be defined by

$$f_i(t_i)[g, a_i] = \langle h_i(t_i|a_i, g) \rangle^\delta$$

for all a_i and $g \in \overline{\mathcal{G}}$. Let $T_i \subseteq \left\{ \widehat{t}_i : \overline{\mathcal{G}} \times A \rightarrow \left\{ 0, \delta, \dots, \langle 2M \rangle^\delta \right\} \right\}$ be the range of f_i . Note that T_i is finite set by construction.

Define $\widehat{\pi}_i : T_i \rightarrow \Delta(T_j \times \Theta)$ as follows. For each $\widehat{t}_i \in T_i$, fix any $t_i \in T^*$ such that $f_i(t_i) = \widehat{t}_i$. Label this type $\zeta_i(\widehat{t}_i)$ and let

$$\widehat{\pi}_i(\widehat{t}_i) [\{(t_j, \theta)\}] = \pi_i^*(\zeta_i(\widehat{t}_i)) [\{(t_j, \theta) : f_j(t_j) = \widehat{t}_j\}].$$

Fix $\varepsilon \in \left\{ 0, \delta, \dots, \langle 2M \rangle^\delta \right\}$. Let $S_i(\widehat{t}_i) = R_i(\zeta_i(\widehat{t}_i), g, \varepsilon)$.

We argue that S is an ε -best response set on the type space $(T_i, \widehat{\pi}_i)_{i=1,2}$. To see why, observe that

$$a_i \in R_i(\zeta_i(\widehat{t}_i), g, \varepsilon)$$

implies that there exists $\nu \in \Delta(T^* \times \Theta \times A_j)$ such that

$$\begin{aligned} \nu[\{(t_j, \theta, a_j) : a_j \in R_j(t_j, g, \varepsilon)\}] &= 1 \\ \text{marg}_{T^* \times \Theta} \nu &= \pi_i^*(\zeta_i(\widehat{t}_i)) \\ \int_{(t_j, \theta, a_j)} \begin{bmatrix} g_i(a_i, a_j, \theta) \\ -g_i(a'_i, a_j, \theta) \end{bmatrix} d\nu &\geq -\varepsilon \text{ for all } a'_i \in A_i \end{aligned}$$

Now define $\widehat{\nu} \in \Delta(T_j \times \Theta \times A_j)$ by

$$\widehat{\nu}(\widehat{t}_j, \theta, a_j) = \nu[\{(t_j, \theta, a_j) : f_j(t_j) = \widehat{t}_j\}]$$

By construction,

$$\begin{aligned} \widehat{\nu}[\{(\widehat{t}_j, \theta, a_j) : a_j \in S_j(\widehat{t}_j)\}] &= 1 \\ \text{marg}_{T_j \times \Theta} \widehat{\nu} &= \widehat{\pi}_i(\widehat{t}_i) \\ \sum_{a_i, a_j, \widehat{t}_j, \theta} \begin{bmatrix} g_i(a_i, a_j, \theta) \\ -g_i(a'_i, a_j, \theta) \end{bmatrix} \widehat{\nu}(\widehat{t}_j, \theta, a_j) &\geq -\varepsilon \text{ for all } a'_i \in A_i \end{aligned}$$

So $a_i \in BR_i(S)(t_i)$.

Since S is an ε -best response set on the type space $(T_i, \hat{\pi}_i)_{i=1,2}$, $S_i(\hat{t}_i) \subseteq R_i(\hat{t}_i, g, \varepsilon)$.

Thus $a_i \in R_i(t_i, g, \varepsilon) \Rightarrow a_i \in R_i(f_i(t_i), g, \varepsilon)$. ■

Now the same finite type space constructed in the proof of lemma 10 can then be shown to approximately replicate the ε -rationalizability properties of all games in the grid of m action games for all ε :

Lemma 11 *Fix any finite collection of m action games $\bar{\mathcal{G}}$ and $\delta > 0$. There exists a finite type space $(T_i, \hat{\pi}_i)_{i=1,2}$ and functions $(f_i)_{i=1,2}$, each $f_i : T^* \rightarrow T_i$, such that $h_i(f_i(t_i) | a_i, g) \leq h_i(t_i | a_i, g) + \delta$ for all $t_i, g \in \bar{\mathcal{G}}$ and a_i .*

And the same finite type space approximately replicates the ε -rationalizability properties of all m action games:

Lemma 12 *Fix the number of actions m and $\xi > 0$. There exists a finite type space $(T_i, \hat{\pi}_i)_{i=1,2}$ and functions $(f_i)_{i=1,2}$, each $f_i : T^* \rightarrow T_i$, such that $h_i(f_i(t_i) | a_i, g) \leq h_i(t_i | a_i, g) + \xi$ for all $t_i, g \in \mathcal{G}^m$ and a_i .*

Now the proof of lemma 6 can be completed as follows. Lemma 12 implies that for any integer m , there exists a finite type \hat{t}_i^m such that

$$h_i(\hat{t}_i^m | a_i, G) \leq h_i(t_i | a_i, G) + \frac{1}{m}$$

for all a_i and $G \in \mathcal{G}^m$. Now fix any m and let

$$\varepsilon^n = \begin{cases} 2M, & \text{if } n \leq m \\ \frac{1}{n}, & \text{if } n > m \end{cases}.$$

Observe that $\varepsilon^n \rightarrow 0$,

$$h_i(\hat{t}_i^n | a_i, G) \leq h_i(t_i | a_i, G) + 2M = h_i(t_i | a_i, G) + \varepsilon^n$$

for all a_i and $G \in \mathcal{G}^m$ if $n \leq m$ and

$$h_i(\hat{t}_i^n | a_i, G) \leq h_i(t_i | a_i, G) + \frac{1}{n} = h_i(t_i | a_i, G) + \varepsilon^n$$

for all a_i and $G \in \mathcal{G}^m$ if $n > m$.

7 Discussion

The key implication of our denseness result is that there are "enough" finite types to approximate general ones. If there is some strategic behavior that arises in some game on the universal type space, then there exists a finite type for which that behavior arises. In this section we discuss some caveats regarding the interpretation of this result.

- First, we show that there is a sense in which the set of finite types is small.
- Then we discuss a topology that is uniform over all games. The denseness result does not hold with such a topology, and hence the same finite type cannot approximate strategic behavior for an infinite type in all games simultaneously. However, the approximation does hold for all games with a bounded number of actions. Moreover, we believe that in cases where arbitrarily large action spaces are of interest, there is usually a natural metric on actions that makes nearby actions similar.
- Then we discuss relaxing the uniform bound on payoffs that we have used throughout the paper.
- Last, we emphasize the caution needed in working with finite types despite our result.

7.1 Is the set of finite types "generic"?

Our denseness result does not imply that the set of finite types is "generic" in the universal type space. While it is not obvious why this question is important from a strategic point of view, we nonetheless briefly report some results showing that the set of finite types is not generic in either of two standard topological senses.

First, a set is sometimes said to be generic if it is open and dense. But the set of finite types is not open. To show this, it is enough to show that the set of infinite types is dense. This implies that the set of infinite types is not closed and so the set of finite types is not open.

We write T_n^* be the collection of all types that exist on finite belief closed subsets of the universal type space where each player has at most n types. The set of finite types is the countable union $T_F = \cup_n T_n^*$. The set of infinite types is the complement of T_F in T^* .

Theorem 4 *If $\#\Theta \geq 2$, infinite types are dense under the product topology and the strategic topology.*

Proof. It is enough to argue that for an arbitrary n and $t^* \in T_n^*$, we can construct a sequence t^k which converges to t^* in the strategic topology (and thus the product topology) such that each $t^k \notin T_F$. Let $T_1 = T_2 = \{1, \dots, n\}$ and $\hat{\pi}_i : \{1, \dots, n\} \rightarrow \Delta(\{1, \dots, n\} \times \Theta)$. Without loss of generality, we can identify $t^* \in T_n^*$ with type 1 of player 1.

The strategy of proof is simply to allow player i to have an additional signal about $T_j \times \Theta$ (which will require an infinite number of types for each player) but let the informativeness of those signals go to zero. Thus we will have a sequence of types not in T_F but converging to t^* in the strategic topology (and thus the product topology).

Let us suppose each player i observes an additional signal $z_i \in \{1, 2, \dots\}$, so $\tilde{T}_i = T_i \times \{1, 2, \dots\}$, with typical element (n_i, z_i) . Fix $\lambda \in (0, 1)$ and for each $k = 1, 2, \dots$, choose $\tilde{\pi}_i^k : \tilde{T}_i \rightarrow \Delta(\tilde{T}_j \times \Theta)$ to satisfy the following two properties:

$$\left| \tilde{\pi}_i^k((n_j, z_j), \theta | (n_i, z_i)) - (1 - \lambda) \lambda^{z_j - 1} \hat{\pi}_i^k(n_j, \theta | n_i) \right| \leq \frac{1}{k} \quad (2)$$

for all $n_j, z_j, \theta, n_i, z_i$; and

$$\sum_{n_j, z_j} \tilde{\pi}_i^k((n_j, z_j), \theta | (n_i, z_i)) \neq \sum_{n_j, z_j} \tilde{\pi}_i^k((n_j, z_j), \theta' | (n'_i, z'_i)) \quad (3)$$

for all $(\theta, n_i, z_i) \neq (\theta', n'_i, z'_i)$.

Let $t^k \in T^*$ be the type in the universal type corresponding to type $(1, 1)$ in the type space $\left(\tilde{T}_i, \tilde{\pi}_i^k\right)_{i=1,2}$.¹⁴

Now (3) implies that each $t^k \notin T_F$.

We will argue that the sequence t^k converges to t^* in the strategic topology. To see why, let $S_i(n_i, z_i) = R_i(n_i, G, \eta)$ (i.e., the set of η -rationalizable actions of type n_i of player i in game G on the original type space). First observe that S is an η -best-response set in game G on the type space $\left(\tilde{T}_i, \tilde{\pi}_i^\infty\right)_{i=1,2}$. This is true because the type space $\left(\tilde{T}_i, \tilde{\pi}_i^\infty\right)_{i=1,2}$ is equivalent to the original type space $(T_i, \hat{\pi}_i)_{i=1,2}$, where it is common knowledge that each player i observes a

¹⁴Strictly speaking we have only defined the universal type space and finite type spaces so far. The definition of a countable type space, the hierarchy of beliefs it induces (and hence the mapping of a type in such a space into a type in the universal type space) are the obvious and trivial modifications of definition 1 and the constructions that follow it.

conditionally independent draw with probabilities $(1 - \lambda) \lambda^{z_i - 1}$ on $\{1, 2, \dots\}$. But now by (2), S is an $\eta + \frac{2M}{k}$ best response set for G . Thus if $a_i \in R_i(t^*, G, \eta)$, then $a_i \in R_i(t^k, G, \eta + \frac{2M}{k})$. Thus the sequence (t^k, t^*) satisfies the lower strategic convergence property. By corollary 2, this implies strategic convergence. By theorem 1, we also have product convergence. ■

Thus the "open and dense" genericity criterion does not discriminate between finite and infinite types. A more demanding topological genericity criterion is that of "first category". A set is first category if it is the countable union of closed sets with empty interiors. Intuitively, a first category set is small or "non-generic". For example, the set of rationals is dense in the interval $[0, 1]$ but not open and not first category.

Theorem 5 *If $\#\Theta \geq 2$, the set of finite types is first category in T^* under the product topology and under the strategic topology.*

Proof. Theorem 4 already established that the closure of the set of infinite types is the whole universal type space. This implies that each T_n^* has empty interior (in the product topology and in the strategic topology). Since the set of finite types is the countable union of the T_n^* , it is then enough to establish that each T_n^* is closed, in the product topology and thus in the strategic topology.

Suppose $t^k \rightarrow^* t$ and $t^k \in T_n^*$ for all k . By Mertens and Zamir (1985), \rightarrow^* corresponds to the weak topology on the compact set T^* , and hence the support of the limit of a sequence of convergent measures with support on n points cannot be larger. Let $S^1(t) \subseteq T^*$ be the support of $\pi^*(t)$ projected onto T^* . For each $j = 2, 3, \dots$, let $S^j(t) \subseteq T^*$ be the union of the supports of $\pi^*(t)$ for all $t' \in S^{j-1}(t)$. We will show by induction that for all j , $S^j(t^*)$ has at most n elements.

Suppose that $S^1(t^*)$ has $n' > n$ elements. Then there exists m such that $\delta_m(t^*)$ has support with n' elements. But $\delta_m(t^k)$ has support with at most n elements, for each k . Thus we cannot have $\delta_m(t^k) \rightarrow \delta_m(t^*)$, a contradiction.

Now suppose that $S^j(t^*)$ has at most n elements for all $j \leq J$. Suppose that $S^{J+1}(t^*)$ has $n' > n$ elements. Then there exists m such that the union of the supports of $\delta_m(t)$, for all $t \in S^J(t^*)$, has n' elements. But $\delta_m(t^k)$ has support with at most n elements, for each k . Thus we cannot have $\delta_m(t^k) \rightarrow \delta_m(t^*)$, a contradiction.

Thus $t \in T_n^*$ and T_n^* is closed in the product topology. ■

Thus the set of finite types is not generic under two standard topological notions of genericity.

Heifetz and Neeman (2004) use the non topological notion of "prevalence" to discuss genericity on the universal type space. Their approach builds in a restriction to common prior types, and it is not clear how to extend their approach to non common prior types. However, it is clear that finite common prior types will be a shy (i.e., non-generic) set in their setting.

7.2 A Strategic Topology that is Uniform on Games

The upper and lower convergence conditions we took as our starting point are not uniform over games. We focussed on a class of games with uniformly bounded payoffs because we believe that the set of all games with arbitrarily large action spaces and arbitrary payoffs is not of significant economic interest: when arbitrarily large action spaces seem appropriate, there is usually a natural metric on actions that makes nearby actions are similar, i.e. a constraint on the set of admissible payoff functions. Despite this, it seems useful to understand how our results would change if we did ask for uniformity over games.

A distance on types that is uniform in games is:

$$d^*(t_i, t'_i) = \sup_{a_i, G} d(t_i, t'_i | a_i, G). \quad (4)$$

This metric yields a topology that is finer than that induced by the metric d , so the topology it is finer than necessary for the upper and lower convergence properties that we took as our goal. The proof of our denseness result does not extend to this metric, and we believe that the result itself fails. One reason for this belief is that we conjecture that convergence in this topology is equivalent to convergence in the following uniform topology on beliefs:

$$d^{**}(t_i, t'_i) = \sup_k \sup_{f \in F_k} |E(f | \pi^*(t_i)) - E(f | \pi^*(t'_i))|, \quad (5)$$

where F_k is the collection of bounded functions mapping $T^* \times \Theta$ that are measurable with respect to k th level beliefs. We sketch an argument. First, suppose that $d^{**}(t_i, t'_i) \leq \varepsilon$. Then in any game, any action that is δ -rationalizable for t_i will be $\delta + 4\varepsilon M$ -rationalizable for t'_i .

This implies $d^*(t_i, t'_i) \leq \delta + 4\varepsilon M$. On the other hand, if $d^*(t_i, t'_i) \geq \varepsilon$, then by lemma 4, we can construct a game G and action a_i such that $d(t_i, t'_i | a_i, G) \geq \frac{\varepsilon}{2}$.

An argument of Morris (2002) implies that finite types are not dense in the uniform topology on beliefs.¹⁵

7.3 Bounded versus Unbounded Payoffs

We have studied topologies on the class of games with uniformly bounded payoffs. If arbitrary payoff functions are allowed, we can always find a game in which any two types will play very differently, so the only topology that makes strategic behavior continuous is the discrete topology. From this perspective, it is interesting to note that full surplus extraction results in mechanism design theory (Cremer and McLean (1985), McAfee and Reny (1992)) rely on payoffs being unbounded. Thus it is not clear to us how the results in this paper can be used to contribute to a debate on the genericity of full surplus extraction results.¹⁶

7.4 Interpreting the denseness result

That any type can be approximated with a finite type provides only limited support for the use of *simple* finite type spaces in applications. The finite types that approximate arbitrary types in the universal type space are quite complex. The approximation result shows that finite types could conceivably capture the richness of the universal type space, and does not of course establish that the use of any particular simple type space is without loss of generality.

In particular, applying notions of genericity to the belief-closed subspace of finite types must be done with care. Standard notions of genericity for such finite spaces will not in general correspond to strategic convergence. Therefore, results regarding strategic interactions that hold on such "generic" subsets of the finite spaces need not be close to the results that would obtain with arbitrary type spaces. For example, our results complement those of Neeman (2004) and Heifetz and Neeman (2004) on the drawbacks of analyzing genericity

¹⁵Morris (2002) shows that finite types are not dense in the topology of uniform convergence of higher order expectations. Convergence in the metric d^{**} implies uniform convergence of higher order expectations.

¹⁶Bergemann and Morris (2001) showed that both the set of full surplus extraction types and the set of non full surplus extraction types are dense in the product topology among finite common prior types, and the same argument would establish that they are dense in the strategic topology identified in this paper. But of course it is trivial that neither set is dense in the discrete topology, which is the "right" topology for the mechanism design problem.

with respect to collections of (in their case, priors over) types where beliefs about Θ determine the entire hierarchy of beliefs, as is done, for instance, in Cremer and McLean (1985), McAfee and Reny (1992), and Jehiel and Moldovanu (2001).

8 Appendix

Lemma 7: *For any integer m and $\varepsilon > 0$, there exists a finite collection of m action games $\bar{\mathcal{G}}$ such that, for every $g \in G^m$, there exists $g' \in \bar{\mathcal{G}}$ such that $D(g, g') \leq \varepsilon$.*

Proof of Lemma 7. For any integer N , we write

$$\mathcal{G}_N = \left\{ g : \{1, \dots, m\}^2 \times \Theta \rightarrow \left\{ -M, -M + \frac{1}{N}, -M + \frac{2}{N}, \dots, M - \frac{1}{N}, M \right\}^2 \right\}.$$

For any game g , choose $g' \in \mathcal{G}_N$ to minimize $D(g, g')$. Clearly $D(g, g') \leq \frac{1}{2N}$. ■

Lemma 8: *For all i, t_i, a_i, g and g' ,*

$$h_i(t_i|a_i, g) \leq h_i(t_i|a_i, g') + 2D(g, g')$$

for all i, a_i .

Proof of Lemma 8. By the definition of R , we know that $R(g', \delta)$ is a δ -best response set for g' . Thus $R(g', \delta)$ is a $(\delta + 2D(g, g'))$ -best response set for g . So $R_i(t_i, g', \delta) \subseteq R_i(t_i, g, \delta + 2D(g, g'))$. Now if $a_i \in R_i(t_i, g', \delta)$, then $a_i \in R_i(t_i, g, \delta + 2D(g, g'))$. So $\delta \geq h_i(t_i|a_i, g')$ implies $\delta + 2D(g, g') \geq h_i(t_i|a_i, g)$. So $h_i(t_i|a_i, g') + 2D(g, g') \geq h_i(t_i|a_i, g)$ ■

Lemma 9: *For any integer m and $\varepsilon > 0$, there exists a finite collection of m action games $\bar{\mathcal{G}}$ such that, for every $g \in G^m$, there exists $g' \in \bar{\mathcal{G}}$ such that*

$$|h_i(t_i|a_i, g) - h_i(t_i|a_i, g')| \leq \varepsilon.$$

for all i, t_i, a_i, g and g' .

Proof of Lemma 9. By lemma 7, we can choose finite collection of games $\bar{\mathcal{G}}$ such that, for every $g \in \mathcal{G}^m$, there exists $g' \in \bar{\mathcal{G}}$ such that $D(g, g') \leq \frac{\varepsilon}{2}$. Lemma 8 now implies that we also have

$$h_i(t_i|a_i, g) \leq h_i(t_i|a_i, g') + \varepsilon$$

and

$$h_i(t_i|a_i, g') \leq h_i(t_i|a_i, g) + \varepsilon.$$

Lemma 11: Fix any finite collection of m action games $\overline{\mathcal{G}}$ and $\delta > 0$. There exists a finite type space $(T_i, \hat{\pi}_i)_{i=1,2}$ and functions $(f_i)_{i=1,2}$, each $f_i : T^* \rightarrow T_i$, such that $h_i(f_i(t_i) | a_i, g) \leq h_i(t_i | a_i, g) + \delta$ for all $t_i, g \in \overline{\mathcal{G}}$ and a_i . ■

Proof of Lemma 11. We will use the type space constructed in the lemma 10, which had the property that

$$R_i(t_i, g, \varepsilon) \subseteq R_i(f_i(t_i), g, \varepsilon) \quad (6)$$

for all $\varepsilon \in \{0, \delta, 2\delta, \dots\}$. By definition,

$$a_i \in R_i(t_i, g, h_i(t_i | a_i, g)).$$

By monotonicity,

$$a_i \in R_i\left(t_i, g, \langle h_i(t_i | a_i, g) \rangle^\delta\right).$$

By (6),

$$R_i\left(t_i, g, \langle h_i(t_i | a_i, g) \rangle^\delta\right) \subseteq R_i\left(f_i(t_i), g, \langle h_i(t_i | a_i, g) \rangle^\delta\right)$$

Thus

$$h_i(f_i(t_i) | a_i, g) \leq \langle h_i(t_i | a_i, g) \rangle^\delta \leq h_i(t_i | a_i, g) + \delta. \quad \blacksquare$$

Lemma 12: Fix the number of actions m and $\xi > 0$. There exists a finite type space $(T_i, \hat{\pi}_i)_{i=1,2}$ and functions $(f_i)_{i=1,2}$, each $f_i : T^* \rightarrow T_i$, such that $h_i(f_i(t_i) | a_i, g) \leq h_i(t_i | a_i, g) + \xi$ for all $t_i, g \in G^m$ and a_i .

Proof of lemma 12. Fix m and $\xi > 0$. By lemma 9, there exists a finite collection of m -action games $\overline{\mathcal{G}}$ such that for every finite-action game g , there exists $g' \in \overline{\mathcal{G}}$ such that

$$h_i(t_i | a_i, g) \leq h_i(t_i | a_i, g') + \frac{\xi}{3} \quad (7)$$

$$\text{and } h_i(t_i | a_i, g') \leq h_i(t_i | a_i, g) + \frac{\xi}{3} \quad (8)$$

for all i, t_i and a_i . By lemma 11, there exists a finite type space $(T_i, \hat{\pi}_i)_{i=1,2}$ and functions $(f_i)_{i=1,2}$, each $f_i : T^* \rightarrow T_i$, such that

$$h_i(f_i(t_i) | a_i, g) \leq h_i(t_i | a_i, g) + \frac{\xi}{3} \quad (9)$$

for all $t_i, g \in \overline{\mathcal{G}}$ and a_i .

Now fix any i, t_i, a_i and g . By (7), there exists g' such that

$$h_i(t_i|a_i, g') - h_i(t_i|a_i, g) \leq \frac{\xi}{3}$$

and

$$h_i(f_i(t_i)|a_i, g) - h_i(f_i(t_i)|a_i, g') \leq \frac{\xi}{3}.$$

By (9),

$$h_i(f_i(t_i)|a_i, g') - h_i(t_i|a_i, g') \leq \frac{\xi}{3}.$$

So

$$\begin{aligned} h_i(f_i(t_i)|a_i, g) - h_i(t_i|a_i, g) &\leq \left\{ \begin{array}{l} (h_i(f_i(t_i)|a_i, g) - h_i(f_i(t_i)|a_i, g')) \\ + (h_i(f_i(t_i)|a_i, g') - h_i(t_i|a_i, g')) \\ + (h_i(t_i|a_i, g') - h_i(t_i|a_i, g)) \end{array} \right\} \\ &\leq \xi. \end{aligned}$$

■

References

- [1] Battigalli, P. and M. Siniscalchi (2003). “Rationalization and Incomplete Information,” *Advances in Theoretical Economics* **3**, Article 3.
- [2] Bergemann, D. and S. Morris (2001). “Robust Mechanism Design,” <http://www.princeton.edu/~smorris/pdfs/robustmechanism2001.pdf>.
- [3] Bergemann, D. and S. Morris (2004). “Robust Mechanism Design,” forthcoming in *Econometrica*.
- [4] Brandenburger, A. and E. Dekel (1987). “Rationalizability and Correlated Equilibria,” *Econometrica* **55**, 1391-1402.
- [5] Brandenburger, A. and E. Dekel (1993). “Hierarchies of Beliefs and Common Knowledge,” *Journal of Economic Theory* **59**, 189-198.

- [6] Cremer, J. and R. McLean (1985). "Optimal Selling Strategies Under Uncertainty for a Discriminating Monopolist when Demands are Interdependent," *Econometrica* **53**, 345-361.
- [7] Dekel, E., D. Fudenberg and S. Morris (2005). "Interim Rationalizability," mimeo.
- [8] Ely, J. and M. Peski (2005). "Hierarchies of Belief and Interim Rationalizability," <http://ssrn.com/abstract=626641>.
- [9] Fudenberg, D. and D.K. Levine (1986). "Limit Games and Limit Equilibria," *Journal of Economic Theory* **38**, 261-279.
- [10] Geanakoplos, J. and H. Polemarchakis (1982). "We Can't Disagree Forever," *Journal of Economic Theory* **28**, 192-200.
- [11] Harsanyi, J. (1967/68). "Games with Incomplete Information Played by Bayesian Players, I-III," *Management Science* **14**, 159-182, 320-334, 486-502.
- [12] Heifetz, A. (1993) "The Bayesian Formulation of Incomplete Information—The Non-Compact Case," *The International Journal of Game Theory* **21**, 329-338.
- [13] Heifetz A., and D. Samet (1998). "Topology-Free Typology of Beliefs," *Journal of Economic Theory* **82**, 324-341.
- [14] Heifetz, A. and Z. Neeman (2004). "On the Generic Infeasibility of Full Surplus Extraction in Mechanism Design," <http://www.bu.edu/econ/Working%20Papers/papers/Zvika%20Neeman/wp2004/FSE9.pdf>.
- [15] Jehiel, P. and B. Moldovanu (2001). "Efficient Design with Interdependent Valuations," *Econometrica* **69**, 1237-59.
- [16] Kajii, A. and S. Morris (1997). "Payoff Continuity in Incomplete Information Games," *Journal of Economic Theory* **82**, 267-276.
- [17] Lipman, B. (2003). "Finite Order Implications of Common Priors," *Econometrica* **71**, 1255-1267.

- [18] McAfee, R. P. and P. R. Reny (1992). "Correlated Information and Mechanism Design," *Econometrica* **60**, 395-422.
- [19] Mertens, J.-F., S. Sorin and S. Zamir (1994). "Repeated Games: Part A Background Material," CORE Discussion Paper #9420.
- [20] Mertens, J.-F. and Zamir, S. (1985). "Formulation of Bayesian Analysis for Games with Incomplete Information," *International Journal of Game Theory* **14**, 1-29.
- [21] Monderer, D. and D. Samet (1996). "Proximity of Information in Games with Common Beliefs," *Mathematics of Operations Research* **21**, 707-725.
- [22] Morris, S. (2002). "Typical Types," <http://www.princeton.edu/~smorris/pdfs/typicaltypes.pdf>.
- [23] Munkres, J. (1975). *Topology*. Prentice Hall.
- [24] Neeman, Z. (2004). "The Relevance of Private Information in Mechanism Design," *Journal of Economic Theory* **117**, 55-77.
- [25] Rubinstein, A. (1989). "The Electronic Mail Game: Strategic Behavior under Almost Common Knowledge," *American Economic Review* **79**, 385-391.
- [26] Weinstein, J. and M. Yildiz (2003). "Finite Order Implications of Any Equilibrium," http://econ-www.mit.edu/faculty/download_pdf.php?id=911.
- [27] Wilson, R. (1987). "Game-Theoretic Analyses of Trading Processes," in *Advances in Economic Theory: Fifth World Congress*, ed. Truman Bewley. Cambridge: Cambridge University Press chapter 2, pp. 33-70.