



Topologies on Types

Citation

Dekel, Eddie, Drew Fudenberg, and Stephen Morris. 2006. Topologies on types. *Theoretical Economics* 1, no. 3: 275–309.

Published version

<http://econtheory.org>

Link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:3160489>

Terms of use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material (LAA), as set forth at

<https://harvardwiki.atlassian.net/wiki/external/NGY5NDE4ZjgzNTc5NDQzMGIzZWZhMGFIOWI2M2EwYTg>

Accessibility

<https://accessibility.huit.harvard.edu/digital-accessibility-policy>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#)

Topologies on types

EDDIE DEKEL

Department of Economics, Northwestern University and School of Economics, Tel Aviv University

DREW FUDENBERG

Department of Economics, Harvard University

STEPHEN MORRIS

Department of Economics, Princeton University

We define and analyze a “strategic topology” on types in the Harsanyi-Mertens-Zamir universal type space, where two types are close if their strategic behavior is similar in all strategic situations. For a fixed game and action define the distance between a pair of types as the difference between the smallest ε for which the action is ε interim correlated rationalizable. We define a *strategic topology* in which a sequence of types converges if and only if this distance tends to zero for any action and game. Thus a sequence of types converges in the strategic topology if that smallest ε does not jump either up or down in the limit. As applied to sequences, the upper-semicontinuity property is equivalent to convergence in the product topology, but the lower-semicontinuity property is a strictly stronger requirement, as shown by the electronic mail game. In the strategic topology, the set of “finite types” (types describable by finite type spaces) is dense but the set of finite common-prior types is not.

KEYWORDS. Rationalizability, incomplete information, common knowledge, universal type space, strategic topology.

JEL CLASSIFICATION. C70, C72.

1. INTRODUCTION

Harsanyi (1967–68) proposed, and Mertens and Zamir (1985) constructed, a universal type space into which (under some technical assumptions) any incomplete information about a strategic situation can be embedded. As a practical matter, applied researchers do not work with that type space but with smaller subsets of the universal type space.

Eddie Dekel: dekel@northwestern.edu

Drew Fudenberg: fudenberg@husc.harvard.edu

Stephen Morris: smorris@princeton.edu

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, the Third International PIER/IGIER Conference on Economic Theory, the 2005 NSF/CEME Mathematical Economics Conference at Berkeley and the Harvard/MIT theory workshop. We are grateful to Dirk Bergemann, Jeff Ely, Yossi Feinberg, Aviad Heifetz, Atsushi Kajii, Bart Lipman, Aldo

Copyright © 2006 Eddie Dekel, Drew Fudenberg, and Stephen Morris. Licensed under the [Creative Commons Attribution-NonCommercial License 2.5](http://creativecommons.org/licenses/by-nc/2.5/). Available at <http://econtheory.org>.

Mertens and Zamir show that finite types 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. Three ingredients need to be formalized in this approach: how we vary the “strategic situations,” what is meant by “strategic behavior” (i.e., what solution concept), and what is meant by “similar.”

To define “strategic situations,” we start with a given space of uncertainty, Θ , and a type space over that space, i.e. all possible beliefs and higher order beliefs about Θ . We then study the effect of changing the action sets and payoff functions while holding the type space fixed. 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 between the type space and the strategic situation is standard in the mechanism-design literature, and it seems necessary for any sort of comparative statics analysis, but it is at odds with the interpretation of the universal type space as describing all possible uncertainty, including uncertainty about the payoff functions and actions. According to this latter view one cannot identify “higher order beliefs” independent of payoffs in the game.² In contrast, our definition of a strategic topology relies crucially on making this distinction.

Our notion of “strategic behavior” is the set of interim-correlated-rationalizable actions that we analyze in Dekel et al. (2006). This set of actions is 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) conjectures about which actions are played at a given type profile and payoff-relevant state. Under interim correlated rationalizability, a player’s conjectures allow for arbitrary correlation between other players’ actions, and 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). A key advantage of this solution concept for our purposes is that all type spaces that have the same hierarchies of beliefs have the same set of interim-correlated-rationalizable outcomes, so it is a solution concept that can be characterized by working with the universal type space.

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 interim-correlated-rationalizable actions

Rustichini, Muhamet Yildiz, and Shmuel Zamir for helpful discussions; Marcin Peški for detailed comments; and Satoru Takahashi for expert proof reading. We thank the referees and an editor for their detailed and helpful comments. This material is based upon work supported by the National Science Foundation under Grants 0426199, 0111830, and 0095321. Morris is grateful for financial support from the John Simon Guggenheim Foundation and the Center for Advanced Studies in Behavioral Economics.

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).

²See the discussion in Mertens et al. (1994, Remark 4.20b).

has the continuity properties that the best-response correspondences, rationalizable actions, and Nash equilibria all have in complete-information games, while still being coarse enough to be useful.³ A review of those properties helps clarify our work. Fix a family of complete-information games with payoff functions 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 hemicontinuous but not lower hemicontinuous 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,⁴ $BR_i(a_j, \varepsilon, \lambda)$, 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. That is, the ε that measures the departure from best response or equilibrium is continuous. The strategic topology is the coarsest metric topology with this continuity property.

Thus for a fixed game and action, we identify for each type of a player, the smallest ε for which the action is ε interim correlated rationalizable. The distance between a pair of types (for a fixed game and action) is the difference between those smallest ε 's. In our strategic topology a sequence converges if and only if this distance tends to zero pointwise for any action and game. Thus a sequence of types converges in the strategic topology if, for any game and action, the smallest ε does not jump either up or down in the limit, so that the map from types to ε is both upper semicontinuous and lower semicontinuous. We show that a sequence has the upper-semicontinuity property if and only if it converges in the product topology (Theorem 2), that a sequence has the lower-semicontinuity property if and only if it converges in the strategic topology, and that if a sequence has the lower-semicontinuity property then it converges in the product topology (Theorem 1).

A version of the electronic mail game shows that the converse is false: a sequence can converge in the product topology but not have the lower-semicontinuity property, so the product topology is strictly coarser than the strategic topology. This has substantive implications. For example, Lipman (2003) shows that finite common-prior types are dense in the product topology, while we show in Section 7.3 that they are not dense in the strategic topology. Similarly, Yildiz (2006)—using the main argument from Weinstein and Yildiz (2003)—shows that types with a unique interim correlated rationalizable action are open and dense in the product topology; but in the strategic topology there are open sets of types with multiple interim correlated rationalizable actions.⁵

³Topology P is finer than topology P' if every open set in P' is 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 other, non-metrizable, topologies have this property.

⁴An action is an ε -best response if it gives a payoff within ε of the best response.

⁵As will become clear, this follows directly from the definition of the strategic topology: for any type

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—defined as the ε -correlated-interim-rationalizable actions—of 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 I 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” in incomplete-information games. These earlier papers defined topologies on priors or partitions in common-prior information systems with a countable number of types, and used equilibrium as the solution concept.⁶ We do not have a characterization of our strategic topology in terms of beliefs, so we are unable to pin down the relation to these earlier papers.

We use interim correlated rationalizability as the benchmark for rational play. There are two reasons for this choice. First, interim correlated rationalizability depends only on hierarchies of beliefs, and hence is suitable for analysis using the universal type space. In contrast, two types with the same hierarchy of beliefs may have different sets of Nash equilibrium strategies and interim-independent-rationalizable strategies: this is because they can correlate their play on payoff-irrelevant signals using what [Mertens and Zamir \(1985\)](#) call redundant types. If one defined a strategic topology with a solution concept that is not determined by hierarchies of beliefs, one would have to decide what to do about the sensitivity of other solution concepts to “redundant types,” i.e., types with the same hierarchy of beliefs but different correlation possibilities. Second, due to our focus on the universal type space, we have chosen not to impose a common prior on the beliefs. In [Dekel et al. \(2006\)](#), we argue that interim correlated rationalizability characterizes the implications of common knowledge of rationality without the common prior assumption. [Dekel et al. \(2004\)](#) argue that the notion of equilibrium without a common prior has neither an epistemic nor a learning-theoretic foundation. In short, we think this is the natural solution concept for this problem.

The paper is organized as follows. [Section 2](#) reviews the electronic mail game and the failure of the lower-semicontinuity property (but not the upper-semicontinuity property) of interim-correlated-rationalizable outcomes with respect to the product topology. The universal type space is described in [Section 3](#) and the incomplete information games and interim-correlated-rationalizable outcomes we 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, including the “genericity” of finite types, the role of the common prior assumption, and an alternative stronger uniform strategic topology on types. All proofs not contained in the body of the paper are provided in the appendix.

with multiple actions that are ε rationalizable for some strictly negative ε (i.e., actions that are “strictly” rationalizable), there is an open set around that point with multiple rationalizable actions (with $\varepsilon = 0$).

⁶For [Monderer and Samet \(1996\)](#), an information system is a collection of partitions on a fixed state space with a given prior. For [Kajii and Morris \(1997\)](#), an information system is a prior on a fixed type space.

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 a lower-semicontinuity property in the product topology (defined formally below). Specifically, we 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 0-rationalizable for $t_{i\infty}$ but is not ε -rationalizable for t_{ik} for any $\varepsilon < \frac{1}{2}$ and finite k . Thus the minimal ε for which the action is ε -rationalizable jumps down in the limit, and the lower-semicontinuity property discussed in the introduction is not satisfied. Intuitively, for interim-correlated-rationalizable play, the tails of higher order beliefs matter, but the product topology is insensitive to the tails.

On the other hand the set of rationalizable actions does satisfy an upper-semicontinuity property with respect to the sequence of types t_{1k} converging in the product topology to $t_{1\infty}$: since every action is 0-rationalizable for type $t_{1\infty}$, the minimum ε cannot jump up in the limit. In Section 6 we show that product convergence is equivalent to this upper-hemicontinuity property in general.

EXAMPLE 1. 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:

$$\theta = 0: \begin{array}{cc} & \begin{array}{cc} N & I \end{array} \\ \begin{array}{c} N \\ I \end{array} & \begin{array}{|cc|} \hline 0,0 & 0,-2 \\ \hline -2,0 & -2,-2 \\ \hline \end{array} \end{array}$$

In payoff state 1, payoffs are given by:

$$\theta = 1: \begin{array}{cc} & \begin{array}{cc} N & I \end{array} \\ \begin{array}{c} N \\ I \end{array} & \begin{array}{|cc|} \hline 0,0 & 0,-2 \\ \hline -2,0 & 1, 1 \\ \hline \end{array} \end{array}$$

Player i ’s types are $T_i = \{t_{i1}, t_{i2}, \dots\} \cup \{t_{i\infty}\}$. Beliefs are generated by the common prior on the type space given in Figure 1, where $\alpha, \delta \in (0, 1)$. There is a 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 \dots ($k - 2$ 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 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 -interim-correlated-rationalizable actions in this game. We provide a formal definition shortly, but the idea is that we iteratively delete an action for a type at round k if that action is not an ε -best response for any conjecture over the action-type pairs of the opponent that survived to round $k - 1$.

		t_{21}	t_{22}	t_{23}	t_{24}	\dots	$t_{2\infty}$
$\theta = 0$:	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

		t_{21}	t_{22}	t_{23}	t_{24}	\dots	$t_{2\infty}$
$\theta = 1$:	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

FIGURE 1. The common prior on the type space.

Clearly, both N and I are 0 interim correlated rationalizable for types $t_{1\infty}$ and $t_{2\infty}$ of players 1 and 2, respectively. But action N is the unique ε -interim-correlated-rationalizable action for all types of each player i except $t_{i\infty}$, for every $\varepsilon < (1 + \alpha)/(2 - \alpha)$ (note that $(1 + \alpha)/(2 - \alpha) > \frac{1}{2}$). To see this fix any $\varepsilon < (1 + \alpha)/(2 - \alpha)$. 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. \diamond

The example shows that strategic outcomes are not continuous in the product topology. Yildiz (2006)—using the main argument from Weinstein and Yildiz (2003)—shows that this discontinuity with respect to the product topology in the email example is quite general: for any type with multiple interim correlated rationalizable actions, there is a sequence of types converging to it in the product topology with a unique interim correlated rationalizable action for every type in the sequence. Thus the denseness of finite types in the product topology does not imply that they are dense in our topology.

3. TYPES

Games of incomplete information are modelled using type spaces. In this paper we work primarily with the “universal type space” developed by [Mertens and Zamir \(1985\)](#).⁷ This type space is called “universal” because it can be used to embed the belief hierarchies (defined below) that are derived from arbitrary type spaces. In this paper, we occasionally construct finite and countable type spaces, but we do not need to work with general uncountable type spaces, so we do not develop the machinery and assumptions needed to handle them. We review the relevant concepts here.

The set of agents is $I = \{1, 2\}$; we also denote them by i and $j = 3 - i$.⁸ Let Θ be a finite set representing possible payoff-relevant moves by Nature.⁹ Throughout the paper, we write $\Delta(\Omega)$ for the set of probability measures on the Borel field of any topological space Ω ; when Ω is finite or countable we use the Borel field corresponding to the discrete topology, so that all subsets of Ω are measurable.

DEFINITION 1. A countable (finite) type space is any collection $(T_i, \pi_i)_{i \in I}$ where T_i is a countable (finite) set and $\pi_i : T_i \rightarrow \Delta(T_j \times \Theta)$.

Let $X_0 = \Theta$, $X_1 = X_0 \times \Delta(X_0)$, and continuing in this way, for each $k \geq 1$, let

$$X_k = X_{k-1} \times \Delta(X_{k-1}),$$

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. Note that because Θ is finite, each X_k and $\Delta(X_k)$ is compact. An element $(\delta_1, \delta_2, \dots) \in \times_{k=0}^{\infty} \Delta(X_k)$ is a hierarchy of beliefs.

Next we show how to calculate the hierarchy of beliefs associated with a given countable type space.

DEFINITION 2. Given a countable type space $(T_i, \pi_i)_{i \in I}$, for each $k = 1, 2, \dots$, define the k^{th} level beliefs for each type as follows. The *first-level beliefs* are $\hat{\pi}_{i,1} : T_i \rightarrow \Delta(X_0)$, where

$$\hat{\pi}_{i,1} [t_i](\theta) = \sum_{t_j \in T_j} \pi_i [t_i](t_j, \theta).$$

Now we define the k^{th} level beliefs $\hat{\pi}_{i,k+1} : T_i \rightarrow \Delta(X_k)$ inductively, noting that $X_k = (\prod_{l=0}^{k-1} \Delta(X_l)) \times \Theta$:

$$\hat{\pi}_{i,k+1} [t_i]((\delta_{j,l})_{l=1}^k, \theta) = \sum_{\{t_j \in T_j : \hat{\pi}_{j,l}(t_j) = \delta_{j,l}, l=1, \dots, k\}} \pi_i [t_i](t_j, \theta).$$

Let $\hat{\pi}_i(t_i) = (\hat{\pi}_{i,k}(t_i))_{k=1}^{\infty}$.

⁷See also [Brandenburger and Dekel \(1993\)](#), [Heifetz \(1993\)](#), and [Mertens et al. \(1994\)](#). 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 et al. \(2006\)](#).

⁸We 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.

⁹We 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 .

Mertens and Zamir (1985) show the existence of a subset of hierarchies, $T^* \subseteq \times_{k=0}^{\infty} \Delta(X_k)$, and a function $\pi^* : T^* \rightarrow \Delta(T^* \times \Theta)$ that preserves beliefs (i.e., $\text{marg}_{X_k} \pi^*(t) = \delta_{k+1}(t)$), that—by defining $T_i^* = T^*$ and $\pi_i^* = \pi^*$ —generate a *universal type space* $(T_i^*, \pi_i^*)_{i \in I}$ into which all “suitably regular” type spaces can be embedded, in the sense that any hierarchy of beliefs generated by such a type space is an element of T^* .¹⁰ In particular, finite and countable type spaces are covered by the Mertens and Zamir result, and the function $\hat{\pi}_i$ constructed above maps any type from a countable type space into T^* . The same mapping sends the set of types in any countable type space into a subset of T_i^* ; this subset is a “belief-closed subspace” of the universal type space.

DEFINITION 3. A *countable belief-closed subspace* is a collection $(T_i, \pi_i^*)_{i \in I}$ where $T_i \subseteq T_i^*$ is countable and $\pi_i^*(t_i)[T_j \times \Theta] = 1$ for all $i \in I$ and all $t_i \in T_i$. It is finite if T_i is finite for $i = 1, 2$.

While we construct and use countable type spaces, we are interested only in the hierarchies of beliefs generated by those type spaces, and hence the belief-closed subspaces into which they are mapped. Therefore, whenever type spaces $(T_i, \pi_i)_{i \in I}$ are constructed, we abuse notation and view the types $t_i \in T_i$ as elements of the universal type space with $\pi_i^*(t_i)$ being the belief over j 's types and Nature, rather than writing the more cumbersome $\pi_i^*(\hat{\pi}_i(t_i))$.

Since our main result is that finite types are dense in the universal type space, we need to define finite types.

DEFINITION 4. A type $t_i \in T_i^*$ is *finite* if it is an element of a finite belief-closed subspace, i.e., there exists $(T_j)_{j \in I}$ such that $t_i \in T_i$ and $(T_j, \pi_j^*)_{j \in I}$ is a finite belief-closed subspace. A type is *infinite* if it is not finite.

REMARK 1. If a type space $(T_i, \pi_i)_{i \in I}$ is finite then each $t_i \in T_i$ corresponds to a finite type in the universal type space. However, an infinite type space can contain some finite types, for two reasons. First, an infinite number of types can be mapped to the same type in the universal type space; for example, a complete-information game where Θ is a singleton and so each player necessarily has a single belief hierarchy can be combined with a publicly-observed randomizing device to create an infinite type space. Second, a finite belief-closed subspace can always be combined with a disjoint infinite type space to yield an infinite type space that contains finite types; the types $\{t_{1\infty}, t_{2\infty}\}$ in the email game are an example of this.

REMARK 2. To test whether a given type in a type space is finite, it must be mapped into its hierarchy of beliefs, i.e., into its image in the universal type space. A given type t in the universal type space is finite if and only if it “reaches” only a finite set of types. We write $r(t) = \text{support}(\text{marg}_{T^*} \pi^*(t))$ for the set of types of the opponent directly reached by type t . If $t' \in r(t)$, we say that t' is reached in one step from t . If type $t'' \in r(t')$

¹⁰The Mertens-Zamir result applies to type spaces that are Polish and given the Borel sigma field. Heifetz and Samet (1998) extend the result to a class of non-topological type spaces.

and $t' \in r(t)$, we say that t'' is reached in two steps from t . And so on. Now if we set $Z(-1, t) = \emptyset, Z(0, t) = \{t\}$, and, for $k \geq 1$,

$$Z(k, t) = Z(k - 2, t_1) \cup \left(\bigcup_{t' \in Z(k-1, t_1)} r(t_1) \right),$$

then, for k odd, $Z(k, t)$ is the set of types of the opponent reached in k or less steps from t ; and, for k even, $Z(k, t)$ is the set of types of the same player as t reached in k or less steps. Now type t is finite if and only if $Z(k, t)$ is bounded above.¹¹

Recall that our intent is to find a strategic topology for the universal type space. However, types in the universal type space are just hierarchies of beliefs, so we consider topologies on hierarchies of beliefs. For this approach to be sensible, it is important that we base our topology on a solution concept that depends only on the hierarchy of beliefs and not on other aspects of the type space. As noted, in our companion paper we show that interim correlated rationalizability has this property.

4. GAMES AND INTERIM CORRELATED RATIONALIZABILITY

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. Note the assumption of a uniform bound on payoffs: If payoffs can be arbitrarily large, then best responses, rationalizable sets, etc. are unboundedly sensitive to beliefs, and as we will see our “strategic topology” reduces to the discrete one; we elaborate on this point in Section 7.6. The topology we define is independent of the value of the payoff bound M so long as M is finite.

Here we restate definitions and results from our companion paper, Dekel et al. (2006). In that paper, we vary the type space and hold fixed the game G being played and characterize 0-rationalizable actions. 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 examine ε -rationalizable actions, so we need to make the dependence of the solution on G and ε explicit. The companion paper defines interim correlated rationalizability on arbitrary type spaces, and shows that two types that have the same hierarchy of beliefs (and so map to the same point in the universal type space T^*) have the same set of ε -interim-correlated-rationalizable actions for any ε . Thus in this paper we can without loss of generality specialize the definitions and results to the type space T^* .

For any subset of actions for all types, we first define the best replies when conjectures are restricted to those actions. Let $\sigma_j : T_j^* \times \Theta \rightarrow \Delta(A_j)$ denote player i 's conjecture about the distribution of the player j 's action as a function of j 's type and the state of Nature. For any measurable σ_j and any belief over opponents' types and the state of Nature, $\pi_i^*(t_i) \in \Delta(T_j^* \times \Theta)$, let $\nu(\pi_i^*(t_i), \sigma_j) \in \Delta(T_j^* \times \Theta \times A_j)$ denote the induced joint conjecture over the space of types, Nature, and actions, where for measurable $F \subset T_j^*$,

$$\nu(\pi_i^*(t_i), \sigma_j)(F \times \{(\theta, a_j)\}) = \int_F \sigma_j(t_j, \theta)[a_j] \cdot \pi_i^*(t_i)[dt_j, \theta].$$

¹¹The choice of topology for defining the support of a set is irrelevant since all we care about here is whether or not it is finite.

DEFINITION 5. Given a specification of a subset of actions for each possible type of opponent, denoted by $E_j = ((E_{t_j})_{t_j \in T_j^*})$, 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

$$BR_i(t_i, E_j, G, \varepsilon) = \left\{ a_i \in A_i : \begin{array}{l} \exists v \in \Delta(T_j^* \times \Theta \times A_j) \text{ such that} \\ (1) v[\{(t_j, \theta, a_j) : a_j \in E_{t_j}\}] = 1 \\ (2) \text{marg}_{T_j^* \times \Theta} v = \pi_i^*(t_i) \\ (3) \int_{T_j^* \times \Theta \times A_j} (g_i(a_i, a_j, \theta) - g_i(a'_i, a_j, \theta)) dv \\ \geq -\varepsilon \text{ for all } a'_i \in A_i \end{array} \right\}.$$

The correspondence of ε best replies in game G for all types given a subset of actions for all types is denoted $BR(G, \varepsilon) : ((2^{A_i})_{i \in I}^{T_i^*}) \rightarrow ((2^{A_i})_{i \in I}^{T_i^*})$ and defined by $BR(G, \varepsilon)(E) = ((BR_i(t_i, E_{-i}, G, \varepsilon))_{t_i \in T_i})_{i \in I}$, where $E = ((E_{t_i})_{t_i \in T_i})_{i \in I} \in ((2^{A_i})_{i \in I}^{T_i^*})$.¹²

REMARK 3. In cases where E_j is not measurable, we interpret $v[\{(t_j, \theta, a_j) : a_j \in E_j\}] = 1$ as saying that there is a measurable subset $E' \subseteq E_j$ such that $v[\Theta \times E'] = 1$. Because $A_{-i} \times \Theta$ is finite and utility depends only on actions and conjectures, the set of ε best responses in G given some E_{-i} , $BR_i(t_i, E_{-i}, G, \varepsilon)$, is non-empty provided there exists at least one measurable σ_{-i} that satisfies (1). Such σ_{-i} exist whenever E_{-i} is non-empty and measurable, and more generally whenever E_{-i} admits a measurable selection.

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

DEFINITION 6. Fix a game $G = (A_i, g_i)_{i \in I}$ and ε .

(i) The *interim-correlated-rationalizable set*,

$$R(G, \varepsilon) \equiv ((R_i(t_i, G, \varepsilon))_{t_i \in T_i^*})_{i \in I} \in ((2^{A_i})_{i \in I}^{T_i^*}),$$

is the largest (in the sense of set inclusion) fixed point of BR .

(ii) The k^{th} -order *interim-correlated-rationalizable sets*, $k = 0, 1, 2, \dots, \infty$, are defined as follows:

$$\begin{aligned} R_0(G, \varepsilon) &\equiv (R_{i,0}(G, \varepsilon))_{i \in I} \equiv ((R_{i,0}(t_i, G, \varepsilon))_{t_i \in T_i^*})_{i \in I} \equiv ((A_i)_{t_i \in T_i^*})_{i \in I} \\ R_k(G, \varepsilon) &\equiv (R_{i,k}(G, \varepsilon))_{i \in I} \equiv ((R_{i,k}(t_i, G, \varepsilon))_{t_i \in T_i^*})_{i \in I} \equiv BR(G, \varepsilon)(R_{k-1}) \\ R_\infty(G, \varepsilon) &\equiv (R_{i,\infty}(G, \varepsilon))_{i \in I} \equiv ((R_{i,\infty}(t_i, G, \varepsilon))_{t_i \in T_i^*})_{i \in I} \equiv \bigcap_{k=1}^\infty R_k(G, \varepsilon). \end{aligned}$$

Dekel et al. (2006) establish that the sets are well-defined, and show the following relationships among them for the case $\varepsilon = 0$; the extensions to general non-negative ε are immediate.

¹²We abuse notation and write BR both for the correspondence specifying best replies for a type and for the correspondence specifying these actions for all types.

RESULT 1. (i) $R(G, \varepsilon)$ equals $R_\infty(G, \varepsilon)$.

(ii) $R_{i,k}(G, \varepsilon)$ and $R_{i,\infty}(G, \varepsilon)$ are measurable functions from $T_i \rightarrow 2^{A_i} / \emptyset$, and for each action a_i and each k the sets $\{t_i : a_i \in R_{i,k}(t_i, G, \varepsilon)\}$ and $\{t_i : a_i \in R_{i,\infty}(t_i, G, \varepsilon)\}$ are closed.

To lighten the paper, we will frequently drop the “interim correlated” modifier, and simply speak of rationalizable sets and rationalizability whenever no confusion will result.

In defining our strategic topology, we exploit the following closure properties of R 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 prove this by induction. It is vacuously true for $k = 0$.

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

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

The sequence $\Psi_{i,k}(t_i, \varepsilon^n)$ is decreasing in n (under set inclusion) and converges (by σ -additivity) to $\Psi_{i,k}(t_i, \varepsilon)$; moreover, in Dekel et al. (2006) we show that each $\Psi_{i,k}(t_i, \varepsilon^n)$ is compact. Let

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

The sequence $\Lambda_{i,k}(t_i, a_i, \varepsilon^n)$ is decreasing in n (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 $a_i \in R_i(t_i, G, \varepsilon)$. Thus for any t_i, a_i , and G ,

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

exists.

PROOF. We have

$$\begin{aligned} a_i &\in R_i(t_i, G, \varepsilon^n) \text{ for all } n \\ &\Rightarrow a_i \in R_{i,k}(t_i, G, \varepsilon^n) \text{ for all } n \text{ and } k \\ &\Rightarrow a_i \in R_{i,k}(t_i, G, \varepsilon) \text{ for all } k, \text{ by Lemma 1} \\ &\Rightarrow a_i \in R_i(t_i, G, \varepsilon). \end{aligned}$$

Since we are considering a fixed finite game, $\inf\{\varepsilon : a_i \in R_i(t_i, G, \varepsilon)\}$ is finite, and the first part of the proposition shows that the infimum is attained. \square

5. THE STRATEGIC TOPOLOGY

5.1 Basic definitions

The most commonly used topology on the universal type space is the “product topology” on the hierarchy.

DEFINITION 7. $t_i^n \rightarrow^* t_i$ if, for each k , $\delta_k(t_i^n) \rightarrow \delta_k(t_i)$ 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.¹³

However, we would like to use a topology that is fine enough that the ε -best-response correspondence and ε -rationalizable sets have the continuity properties satisfied by the ε -best-response correspondence, ε -Nash equilibrium, and ε -rationalizability in complete information games with respect to the payoff functions. The electronic mail game shows that the product topology is too coarse for these continuity properties to obtain, which suggests the use of a finer topology. One way of phrasing our question is whether there is any topology that is fine enough for the desired continuity properties and yet coarse enough that finite types are dense; our main result is that indeed there is: this is true for the “strategic topology” that we are about to define.

Ideally, it would be nice to know that our strategic topology is the coarsest one with the desired continuity properties, but since non-metrizable topologies are hard to analyze, we have chosen to work with a metric topology. Hence we construct the coarsest metric topology with the desired properties.

For any fixed game and feasible 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 for any $G = (A_i, g_i)_{i \in I}$ and $a_i \in A_i$,

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

In extending this to a distance over types, we allow for a larger difference in the h 's in games with more actions. Thus, the metric that we define is not uniform over the number of actions in the game. When studying a game with a large or unbounded number

¹³It is used not only in the constructions by Mertens and Zamir (1985) and Brandenburger and Dekel (1993) but also in more recent work by Lipman (2003) and Weinstein and Yildiz (2003).

of actions, we think there should be a metric on actions and accompanying constraints on the set of admissible payoff functions (such as continuity or single-peakedness) that make “nearby” actions “similar.” Requiring uniformity here would require that strategic convergence be uniform over the number of actions in the game, and this seems too strong a requirement given that we allow arbitrary (bounded) payoff functions.

For any integer m , there is no loss in generality in taking the action spaces to be $A_1^m = A_2^m = \{1, 2, \dots, m\}$. Having fixed the action sets, a game is parameterized by the payoff function g . So for a fixed m , we write g for the game $G = (\{1, 2, \dots, m\}, \{1, 2, \dots, m\}, g)$, and \mathcal{G}^m for the set of all such games g .

Now consider the following notion of distance between types:

$$d(t_i, t'_i) = \sum_m \beta^m \sup_{a_i \in A_i^m, g \in \mathcal{G}^m} d_i(t_i, t'_i | a_i, g),$$

where $0 < \beta < 1$.¹⁴

LEMMA 2. *The distance d is a pseudo-metric.*

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

$$\begin{aligned} d(t_i, t''_i) &= \sum_m \beta^m \sup_{a_i \in A_i^m, g \in \mathcal{G}^m} d(t_i, t''_i | a_i, g) \\ &\leq \sum_m \beta^m \sup_{a_i \in A_i^m, g \in \mathcal{G}^m} (d(t_i, t'_i | a_i, g) + d(t'_i, t''_i | a_i, g)) \\ &\leq \sum_m \beta^m \left(\sup_{a_i \in A_i^m, g \in \mathcal{G}^m} d(t_i, t'_i | a_i, g) + \sup_{a_i \in A_i^m, g \in \mathcal{G}^m} d(t'_i, t''_i | a_i, g) \right) \\ &= \sum_m \beta^m \sup_{a_i \in A_i^m, g \in \mathcal{G}^m} d(t_i, t'_i | a_i, g) + \sum_m \beta^m \sup_{a_i \in A_i^m, g \in \mathcal{G}^m} d(t'_i, t''_i | a_i, g) \\ &= d(t_i, t'_i) + d(t'_i, t''_i). \quad \square \end{aligned}$$

Theorem 1 below implies that $d(t_i, t'_i) = 0 \Rightarrow t_i = t'_i$, so that d is in fact a metric.

DEFINITION 8. The *strategic topology* is the topology generated by d .

¹⁴We do not require a supremum over i in this definition as we can instead consider a game with the indices switched on the action spaces and payoff functions (while we do not restrict attention to symmetric games, the set of all games is symmetric).

To analyze and explain the strategic topology, we characterize its convergent sequences using the following two conditions.¹⁵

DEFINITION 9 (Strategic Convergence).

- (i) $((t_i^n)_{n=1}^\infty, t_i)$ satisfy the *upper strategic convergence property* (written $t_i^n \rightarrow_U t_i$) if for every m , $a_i \in A_i^m$, and $g \in \mathcal{G}^m$, $\limsup_n h_i(t_i^n | a_i, g) \leq h_i(t_i | a_i, g)$. We refer to this property as *upper semicontinuity in n* .
- (ii) $((t_i^n)_{n=1}^\infty, t_i)$ satisfy the *lower strategic convergence property* (written $t_i^n \rightarrow_L t_i$) if for every m , $a_i \in A_i^m$, and $g \in \mathcal{G}^m$, $\liminf_n h_i(t_i^n | a_i, g) \geq h_i(t_i | a_i, g)$. We refer to this property as *lower semicontinuity in n* .

Not that by definition if $t_i^n \rightarrow_U t_i$ then for each m , $g \in \mathcal{G}^m$, and $a_i \in A_i^m$ there exists $\varepsilon^n \rightarrow 0$ such that $h_i(t_i | a_i, g) < h_i(t_i^n | a_i, g) + \varepsilon^n$. The statement for lower semicontinuity is analogous (switching t_i^n and t_i in the implication). **Lemma 11** in the appendix states that for each m the sequence ε^n can be chosen independently of g , so that upper and lower strategic convergence have the following stronger implications: for each m ,

$$\begin{aligned} t_i^n \rightarrow_U t_i &\Rightarrow \exists \varepsilon^n \rightarrow 0 \text{ s.t. } h_i(t_i | a_i, g) < h_i(t_i^n | a_i, g) + \varepsilon^n \\ t_i^n \rightarrow_L t_i &\Rightarrow \exists \varepsilon^n \rightarrow 0 \text{ s.t. } h_i(t_i^n | a_i, g) < h_i(t_i | a_i, g) + \varepsilon^n \end{aligned}$$

for every n , $a_i \in A_i^m$, and $g \in \mathcal{G}^m$.¹⁶

We do not require convergence uniformly over all games, as 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 7.5**.

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

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

¹⁵In metric spaces convergence and continuity can be assessed by looking at sequences (Munkres 1975, p. 190). Since we show in **Section 6.1** that each convergence condition coincides with convergence according to a metric topology (the product and strategic topologies respectively) we conclude that the open sets defined directly from the convergence notions below do define these topologies. We do not know if there are non-metrizable topologies with the same convergent sequences.

¹⁶The upper-strategic-convergence property implies a property that resembles upper hemicontinuity of the interim correlated rationalizability correspondence; specifically, $t_i^n \rightarrow_U t_i$ implies that $a_i \in R_i(t_i^n, g, \varepsilon) \forall n \implies a_i \in R_i(t_i, g, \varepsilon)$. Moreover, in **Theorem 1** below we show that $t_i^n \rightarrow_U t_i$ if and only if $t_i^n \rightarrow^* t_i$, which, together with the preceding observation, implies that $R_i(t_i, g, \varepsilon)$ is upper hemicontinuous in t_i w.r.t. the product topology. In contrast, the lower-strategic-convergence property is weaker than a lower-hemicontinuity-like property, and $R_i(t_i, g, \varepsilon)$ is not lower hemicontinuous in t_i w.r.t. the product topology. Instead, $t_i^n \rightarrow^* t_i$ and $a_i \in R_i(t_i, g, \varepsilon) \implies \exists \varepsilon^n \rightarrow \varepsilon$ with $a_i \in R_i(t_i^n, g, \varepsilon^n) \forall n$. As discussed in the introduction, this is what we should expect by analogy with the continuity properties of solution concepts under complete information.

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

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

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 $\underline{\varepsilon}^n(m) \rightarrow 0$ such that for all $a_i, g \in \mathcal{G}^m$,

$$h_i(t_i | a_i, g) < h_i(t_i^n | a_i, g) + \bar{\varepsilon}^n(m)$$

and $h_i(t_i^n | a_i, g) < h_i(t_i | a_i, g) + \underline{\varepsilon}^n(m).$

Thus

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

□

The strategic topology is thus a metric topology where sequences converge if and only if they satisfy upper and lower strategic convergence, and so the strategic topology is the coarsest metric topology with the desired continuity properties. We now illustrate the strategic topology with the example discussed earlier.

5.2 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 semicontinuity noted in Section 2, while $d(t_{1k}, t_{1\infty}) \not\rightarrow 0$, corresponding to the failure of lower semicontinuity.

The beliefs for the types in that example are as follows.

$$\begin{aligned} \pi_1^*(t_{11})[(t_2, \theta)] &= \begin{cases} 1 & \text{if } (t_2, \theta) = (t_{21}, 0) \\ 0 & \text{otherwise} \end{cases} \\ \pi_1^*(t_{1m})[(t_2, \theta)] &= \begin{cases} \frac{1}{2-\alpha} & \text{if } (t_2, \theta) = (t_{2,m-1}, 1) \\ \frac{1-\alpha}{2-\alpha} & \text{if } (t_2, \theta) = (t_{2m}, 1) \\ 0 & \text{otherwise} \end{cases} \quad \text{for all } m = 2, 3, \dots \\ \pi_1^*(t_{1\infty})[(t_2, \theta)] &= \begin{cases} 1 & \text{if } (t_2, \theta) = (t_{2\infty}, 1) \\ 0 & \text{otherwise} \end{cases} \end{aligned}$$

$$\pi_2^*(t_{21})[(t_1, \theta)] = \begin{cases} \frac{1}{2-\alpha} & \text{if } (t_1, \theta) = (t_{11}, 0) \\ \frac{1-\alpha}{2-\alpha} & \text{if } (t_1, \theta) = (t_{12}, 1) \\ 0 & \text{otherwise} \end{cases}$$

$$\pi_2^*(t_{2m})[(t_1, \theta)] = \begin{cases} \frac{1}{2-\alpha} & \text{if } (t_1, \theta) = (t_{1m}, 1) \\ \frac{1-\alpha}{2-\alpha} & \text{if } (t_1, \theta) = (t_{1,m+1}, 1) \\ 0 & \text{otherwise} \end{cases} \quad \text{for all } m = 2, 3, \dots$$

$$\pi_2^*(t_{2\infty})[(t_1, \theta)] = \begin{cases} 1 & \text{if } (t_1, \theta) = (t_{1\infty}, 1) \\ 0 & \text{otherwise.} \end{cases}$$

One can verify that $t_{1k} \rightarrow^* t_{1\infty}$. This is easiest to see by looking at the table presenting the beliefs: t_{1m} assigns probability 1 to each of the following: $\theta = 1$; 2 assigns probability 1 to $\theta = 1$; 2 assigns probability 1 to 1 assigning probability 1 to $\theta = 1$; and so on up to iterations of length $m - 1$. As $m \rightarrow \infty$, the k^{th} level beliefs converge to those of $t_{1\infty}$ where it is common knowledge that $\theta = 1$.

Note that the only finite types here are $t_{i\infty}$.

Denote by \widehat{G} the e-mail game described earlier. For any $\varepsilon < (1 + \alpha)/(2 - \alpha)$,

$$R_{1,0}(t_1, \widehat{G}, \varepsilon) = \{N, I\} \text{ for all } t_1$$

$$R_{2,0}(t_2, \widehat{G}, \varepsilon) = \{N, I\} \text{ for all } t_2$$

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

$$R_{2,1}(t_2, \widehat{G}, \varepsilon) = \{N, I\}$$

$$R_{1,2}(t_1, \widehat{G}, \varepsilon) = \begin{cases} \{N\} & \text{if } t_1 = t_{11} \\ \{N, I\} & \text{if } t_1 \in \{t_{12}, t_{13}, \dots\} \cup \{t_{1\infty}\} \end{cases}$$

$$R_{2,2}(t_2, \widehat{G}, \varepsilon) = \begin{cases} \{N\} & \text{if } t_2 = t_{21} \\ \{N, I\} & \text{if } t_2 \in \{t_{22}, t_{23}, \dots\} \cup \{t_{2\infty}\} \end{cases}$$

$$R_{1,3}(t_1, \widehat{G}, \varepsilon) = \begin{cases} \{N\} & \text{if } t_1 \in \{t_{11}, t_{12}\} \\ \{N, I\} & \text{if } t_1 \in \{t_{13}, t_{14}, \dots\} \cup \{t_{1\infty}\} \end{cases}$$

$$R_{2,3}(t_2, \widehat{G}, \varepsilon) = \begin{cases} \{N\} & \text{if } t_2 = t_{21} \\ \{N, I\} & \text{if } t_2 \in \{t_{22}, t_{23}, \dots\} \cup \{t_{2\infty}\} \end{cases}$$

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

for $m = 2, 3, \dots$

and

$$\begin{aligned}
 R_{1,2m+1}(t_1, \widehat{G}, \varepsilon) &= \begin{cases} \{N\} & \text{if } t_1 \in \{t_{11}, \dots, t_{1,m+1}\} \\ \{N, I\} & \text{if } t_1 \in \{t_{1,m+2}, t_{1,m+3}, \dots\} \cup \{t_{1\infty}\} \end{cases} \\
 R_{2,2m+1}(t_2, \widehat{G}, \varepsilon) &= \begin{cases} \{N\} & \text{if } t_2 \in \{t_{21}, \dots, t_{2m}\} \\ \{N, I\} & \text{if } t_2 \in \{t_{2,m+1}, t_{2,m+2}, \dots\} \cup \{t_{2\infty}\} \end{cases}
 \end{aligned}$$

for $m = 2, 3, \dots$

So

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

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 $d(t_{1k}, t_{1\infty}) \geq \beta^2(1+\alpha)/(2-\alpha)$ 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 \widetilde{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 $A^\gamma = \{x \in X : \inf_{y \in A} |x - y| \leq \gamma\}$. Then the Prokhorov distance between measures δ and δ' is $d^P(\delta, \delta') = \inf\{\gamma \mid \delta(A) \leq \delta'(A^\gamma) + \gamma \text{ for all } A \in \mathcal{F}\}$, and $\tilde{d}^k(t_i, t'_i) = d^P(\delta_k(t_i), \delta_k(t'_i))$.

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

PROOF. To prove this we 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 playing an action that is rationalizable for the other type.

To define the finite games we 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 $A^1 = \{a \in \mathbb{R}^{|\Theta|} : a_\theta = j/z^1 \text{ for some integer } j, 1 \leq j \leq z^1, \sum_\theta a_\theta = 1\}$; 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 be¹⁷

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

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 allows us to determine the rationalizable sets recursively, starting from the first-order actions and working up.

¹⁷The payoff function given in the text is independent of the payoff bound M , and need not satisfy it if M is small—in that case we can simply multiply the payoff function by a sufficiently small positive number.

Define $l(t_i, a_i^1) = E_{\delta_1(t_i)}[2a_i^1(\theta) - \sum_{\theta'} (a_i^1(\theta'))^2]$. This is the loss to type t_i of choosing a_i^1 when $\delta_1(t_i)$ is a feasible first-order action. For all $c > 0$, there is $b > 0$ such that if $\tilde{d}(\delta_1(t_i), \delta_1(t'_i)) > c$, then $l(t_i, \delta_1(t'_i)) > b$. In the game with a given finite grid, $\delta_1(t_i)$ is not in general feasible, and the rationalizable first-order action(s) for type t_i is the point or points $a_1^*(\delta_1(t_i)) \in A^1$ closest to $\delta_1(t_i)$; picking any other point involves a greater loss. Thus the loss to t_i from playing an element of $a_1^*(\delta_1(t'_i))$ instead of an element of $a_1^*(\delta_1(t_i))$ is at least $l(t_i, \delta_1(t'_i)) - \varepsilon_2 - \varepsilon_3$, where ε_2 is the loss from playing $a_1^*(\delta_1(t_i))$ instead of the (infeasible) $\delta_1(t_i)$, and ε_3 is the absolute value of the difference in i 's payoff from playing $a_1^*(\delta_1(t'_i))$ instead of $\delta_1(t'_i)$. Both ε_2 and ε_3 go to 0 in z^1 , uniformly in t_i , so 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

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

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 any probability measure δ_2^* on D^1 such that for any $X \subset \Theta \times A_1$, $\delta_2^*(X) \leq \delta_2(\{(\theta, \delta_1) : \{\theta\} \times a_1^*(\delta_1) \subset X\})$. 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^* . We extend the domain of δ_2^* to all of $\Theta \times \Delta(\Theta)$ by setting $\delta_2^*(\Theta \times Y) = \delta_2^*(\Theta \times (A^1 \cap Y))$.

Next we claim that if there is a $c > 0$ such that $\tilde{d}^2(t_i, t'_i) > c$, then $\delta_2^*(t_i) \neq \delta_2^*(t'_i)$ for all sufficiently fine grids A^1 on $\Delta(\Theta^1)$. 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^1 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 $\tilde{d}^2(t_i, t'_i) > c$.

As with the case of first-order beliefs and actions, this implies that when $\tilde{d}^2(t_i, t'_i) > c$ there is a \bar{z}^2 and $\varepsilon_2 > 0$ such that

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

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 all n in some subsequence

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

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

$$\exists a_i \in A_i^m, g \in \mathcal{G}^m \text{ s.t. } h_i(t_i | a_i, g) + \varepsilon < h_i(t_i^n | a_i, g) \text{ and} \quad (1)$$

$$\exists a_i \in A_i^m, g \in \mathcal{G}^m \text{ s.t. } h_i(t_i^n | a_i, g) + \varepsilon < h_i(t_i | a_i, g). \quad (2)$$

Now suppose that the lower convergence property holds. Therefore

$$\exists \eta^n \rightarrow 0 \text{ s.t. } h_i(t_i^n | a_i, g) < h_i(t_i | a_i, g) + \eta^n, a_i \in A_i^m, g \in \mathcal{G}^m.$$

This combined with (1) gives a contradiction.

Similarly, upper convergence implies that

$$\exists \eta^n \rightarrow 0 \text{ s.t. } h_i(t_i | a_i, g) < h_i(t_i^n | a_i, g) + \eta^n, a_i \in A_i^m, g \in \mathcal{G}^m.$$

This gives a contradiction when combined with (2). \square

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 are m , $a_i \in A_i^m$, and $g \in \mathcal{G}^m$ such that 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 $v^n \in \Delta(T_j^* \times \Theta \times A_j)$ such that

$$(1) v^n \left[(t_j, \theta, a_j) : a_j \in R_j(t_j, g, h_i(t_i | a_i, g) - \delta) \right] = 1$$

$$(2) \text{marg}_{T^* \times \Theta} v^n = \pi_i^*(t_i^n)$$

$$(3) \int_{(t_j, \theta, a_j)} [g_i(a_i, a_j, \theta) - g_i(a_i', a_j, \theta)] dv^n \geq -h_i(t_i | a_i, g) + \delta \text{ for all } a_i' \in A_i^m.$$

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 (1), (2) and (3) hold for every n and $\nu = \lim_n \nu^n$, we have

$$\begin{aligned} (1^*) \nu \left[\left\{ (t_j, \theta, a_j) : a_j \in R_j(t_j, g, h_i(t_i | a_i, g) - \delta) \right\} \right] &= 1 \\ (2^*) \text{marg}_{T^* \times \Theta} \nu &= \pi_i^*(t_i) \\ (3^*) \int_{(t_j, \theta, a_j)} \left[g_i(a_i, a_j, \theta) - g_i(a'_i, a_j, \theta) \right] d\nu &\geq -h_i(t_i | a_i, g) + \delta \text{ for all } a'_i \in A_i^m. \end{aligned}$$

Here (1*) 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. To see why (2*) holds, note that $\text{marg}_{T^* \times \Theta} \nu^n \rightarrow \text{marg}_{T^* \times \Theta} \nu$, since $\nu^n \rightarrow \nu$. It remains to show that $\pi_i^*(t_i^n) \rightarrow \pi_i^*(t_i)$, i.e., $t_i = \pi_i^{*-1}(\lim \pi_i^*(t_i^n))$, whenever $t_i^n \rightarrow^* t_i$. This can be inferred from the Mertens-Zamir homeomorphism and standard results about the continuity of marginal distributions in the joint, but a direct proof is about as short: Recall that $\delta_k(t)$ is the k^{th} level belief of type t , that $\text{marg}_{X_{k-1}} \pi^*(t) = \delta_k(t)$, and that (by definition) $t_i = \pi_i^{*-1}(\lim \pi_i^*(t_i^n))$ if and only if for all k , $\delta_k(t_i) = \delta_k(\pi_i^{*-1}(\lim \pi_i^*(t_i^n)))$. Now $\delta_k(\pi_i^{*-1}(\lim \pi_i^*(t_i^n))) = \text{marg}_{X_{k-1}} \lim \pi_i^*(t_i^n) = \lim \text{marg}_{X_{k-1}} \pi_i^*(t_i^n) = \lim \delta_k(t_i^n) = \delta_k(t_i)$ by product convergence. This proves (2*); (3*) follows from $\nu^n \rightarrow \nu$.

This implies $a_i \in R_i(t_i, g, h_i(t_i | a_i, g) - \delta)$, a contradiction. □

COROLLARY 1. *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 1**, 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 and show that any game has ε -rationalizable actions that are close to the ε -rationalizable actions of some game in the finite grid. This allows us to work with such finite grids. We also take a finite grid of ε 's, $\{\varepsilon_j\}_{j=1}^m$. We then define maps f_i taking each type of i into a function that specifies for every action of i and every game from the finite grid of games the minimal ε_j under which the action is ε_j rationalizable. Each such function is one of finitely many types for i . We then define a belief hierarchy for each type in this finite set of types by arbitrary taking the belief hierarchy of one of the types in the universal type space to which it is mapped. This gives us a finite type space. We show that this map “preserves ε ”: $a_i \in R_i(t_i, g, \varepsilon) \Rightarrow a_i \in R_i(f_i(t_i), g, \varepsilon)$. 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 interim correlated 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.¹⁸

A distinctive feature of our construction is that we identify types in our constructed type space with sets of ε -rationalizable actions for a finite set of ε 's and a finite set of games. The recent paper of [Ely and Pęski \(2006\)](#) similarly identifies types with sets of rationalizable actions, although for their different purpose (constructing a universal type space for the interim-independent-rationalizability solution concept), no approximation is required.

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

PROOF. The two critical stages in the proof are as follows. We first prove that there is a finite set of m -action games, $\overline{\mathcal{G}}^m$, that approximate the set \mathcal{G}^m .

LEMMA 7. *For any integer m and $\varepsilon > 0$, there exists a finite collection of m action games $\overline{\mathcal{G}}^m$ such that, for every $g \in \mathcal{G}^m$, there exists $g' \in \overline{\mathcal{G}}^m$ such that for all i, t_i and $a_i \in A_i^m$,*

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

Next we use this to prove that there is a finite approximating type space.

LEMMA 8. *Fix the number of actions m and $\xi > 0$. There exists a finite type space $(\tilde{T}_i, \tilde{\pi}_i)_{i=1,2}$ and functions $(f_i)_{i=1,2}$, each $f_i : T^* \rightarrow \tilde{T}_i$, 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 .*

The key step in this proof is constructing the type space, so we present that here. The remaining details are provided in the appendix.

Fix a finite set of games $\overline{\mathcal{G}}^m$. Write $\langle x \rangle^\delta$ for the smallest number in the set $\{0, \delta, 2\delta, \dots\}$ greater than x , and let Γ^m be the set of all maps from $A_i^m \times \overline{\mathcal{G}}^m$ into $\{0, \delta, \dots, \langle 2M \rangle^\delta\}$. We build the type spaces $(\tilde{T}_i, \tilde{\pi}_i)_{i=1,2}$ using subsets of Γ^m as the types. Specifically, define the function $f_i : T_i^* \rightarrow \Gamma^m$ by $f_i(t_i) = (\langle h_i(t_i | a_i, g) \rangle^\delta)_{a_i, g}$.

Let \tilde{T}_i be the range of f_i ; note that \tilde{T}_i is a finite set. Thus for given δ each type of i in the universal type space is mapped into a function that specifies for each one of

¹⁸It is not clear how one could develop a topology based on the equilibrium distribution of play in a setting without a common prior.

the finitely many games and actions of i the smallest multiple j of δ under which that action is $j\delta$ -rationalizable. These functions constitute the types in a finite type space. The beliefs in this finite type space are defined next.

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

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

Now the proof of Lemma 6 can be completed as follows. Lemma 8 implies that for any fixed m and t_i there exists a finite type $\tilde{t}_i^{n,m}$ such that

$$h_i(\tilde{t}_i^{n,m} | a_i, g) \leq h_i(t_i | a_i, g) + \frac{1}{n}$$

for all $a_i \in A_i^m$ and $g \in \mathcal{G}^m$. Thus $\tilde{t}_i^{n,m} \rightarrow_L t_i$ as $n \rightarrow \infty$. □

7. DISCUSSION

7.1 Outline of issues

The key implication of our denseness result is that there are “enough” finite types to approximate general ones. In this section we discuss some caveats regarding the interpretation of this result.

- We show that there is a sense in which the set of finite types is small.
- We show that finite common-prior types are not dense in the set of finite types and thus in the set of all types.
- We discuss approaches to generalizing our results to alternative solution concepts.
- We describe 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.
- We discuss relaxing the uniform bound on payoffs that we have used throughout the paper.
- Finally, we emphasize that caution is needed in working with finite types despite our result.

7.2 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.

Let 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.*

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.

By **Mertens and Zamir (1985)**, \rightarrow^* corresponds to the weak topology on the compact set T^* . We will repeatedly use the following implications of weak convergence. If a sequence of measures μ^k on a metric space X converges weakly to μ , and the support of every μ^k has n or fewer elements, then (a) the support of μ has at most n elements; (b) every element \bar{x} of the support of μ is the limit of a sequence of elements $x^k \in \text{support}(\mu^k)$; moreover, (c) there exists an integer K and, for each $k > K$ and $x \in \text{support}(\mu)$, $\chi_k(x) \in \text{support}(\mu^k)$, such that, (i) for all $x \in \text{support}(\mu)$, $\chi_k(x)$ converges to x and (ii) $x, x' \in \text{support}(\mu)$ and $x \neq x'$ implies $\chi_k(x) \neq \chi_k(x')$ for all $k > K$.

Now recall from **Remark 2** that, for k even, we write $Z(k, t_1)$ for the set of types of player 1 reached in k or less steps from t_1 of player 1; and, for k odd, we write $Z(k, t_1)$ for the set of types of player 2 reached in k or less steps from t_1 . This implies that for even k ,

$$Z(k, t_1) = \{t_1\} \cup (\cup_{t \in Z(k-1, t_1)} r(t))$$

and for odd k ,

$$Z(k, t_1) = (\cup_{t \in Z(k-1, t_1)} r(t)).$$

Now fix an n , and suppose $\bar{t}_1^k \rightarrow^* \bar{t}_1$ and $\bar{t}_1^k \in T_n^*$ for all k . So $\bar{t}_1^k \in T_1^k$, where $T_1^k \times T_2^k$ is a belief-closed type space with $\#T_i^k \leq n$. We will establish inductively the following claim.

CLAIM. Fix any positive integer L . (1) $Z(L, \bar{t}_1)$ has at most n elements. (2) There exists K_L such that for every $k > K_L$ and every $t \in Z(L, t_1)$, there exists $\tau^k(t) \in T_i^k$ such that $\tau^k(t) \rightarrow t$ and $\tau^k(t) \neq \tau^k(t')$ for all $t, t' \in Z(L, t_1)$ with $t \neq t'$ (where $i = 1$ if L is even and $i = 2$ if L is odd).

We first establish the claim for $L = 1$. Since $\pi^*(t^k)$ is a sequence of measures converging to $\pi^*(t)$, we have that (a) $Z(1, \bar{t}_1)$ has at most n elements; (b) there is a sequence $\tau_2^k(t_2) \in T_2^k$ s.t. $\tau_2^k(t_2) \rightarrow t_2$; and (c) there exists K_1 such that for all $k > K_1$, $\tau_2^k(t_2) \neq \tau_2^k(t'_2)$ if $t_2 \neq t'_2$.

Now suppose that the claim holds for all $L \leq \bar{L} - 1$, where \bar{L} is even. We establish the claim for \bar{L} . For $k > K_{\bar{L}-1}$, we know that

$$|\{\bar{t}_1^k\} \cup (\cup_{t_2 \in Z(\bar{L}-1, \bar{t}_1)} r(\tau_2^k(t_2)))| \leq |T_1^k| \leq n.$$

Now $|\{\bar{t}_1^k\} \cup (\cup_{t_2 \in Z(\bar{L}-1, \bar{t}_1)} r(\tau_2^k(t_2)))| \leq n$ implies $|\{\bar{t}_1\} \cup (\cup_{t_2 \in Z(\bar{L}-1, \bar{t}_1)} r(t_2))| \leq n$ since $\bar{t}_1^k \rightarrow \bar{t}_1$, $\tau_2^k(t_2) \rightarrow t_2$ and supports cannot grow. Thus $Z(\bar{L}, \bar{t}_1)$ has at most n elements. Also observe that for each $t_1 \in Z(\bar{L}, \bar{t}_1)$, there is a sequence $\tau_1^k(t_1) \in T_1^k$ s.t. $\tau_1^k(t_1) \rightarrow t_1$. This is true by assumption if $t_1 = \bar{t}_1$; otherwise t_1 is the limit of a sequence of types in $r(t_2^k)$ for some $t_2^k \in T_2^k$. Since $Z(\bar{L}, \bar{t}_1)$ is finite, there exists $K_{\bar{L}}$ such that for all $k > K_{\bar{L}}$ and all $t_1 \in Z(\bar{L}, \bar{t}_1)$, there is a sequence of $\tau_1^k(t_1) \in T_1^k$ converging to t_1 with $\tau_1^k(t_1) \neq \tau_1^k(t'_1)$ if $t_1 \neq t'_1$.

Now suppose that the claim holds for all $L \leq \bar{L} - 1$, where \bar{L} is odd. Essentially the same argument establishes the claim for \bar{L} (apart from labelling, the only difference from the case of even \bar{L} is that we do not have \bar{t}_1 in $Z(k, \bar{t}_1)$).

We have now established the claim for all L by induction, and thus that type \bar{t}_1 is an element of T_n^* . □

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

Heifetz and Neeman (2006) 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. They also show that the generic set has the property that any convex combination of an element in the set with an element of its complement is in the set. This property is satisfied here as well: convex combinations of finite types with infinite types are infinite.

7.3 Types with a common prior

We say that t is a finite common-prior type if it belongs to a finite belief-closed subspace $(T_i, \pi_i^*)_{i=1}^2$ and there is a probability distribution π^* on T^* (the common prior) that assigns positive probability to every type of every player, such that $\pi_i^*(t_i) = \pi^*(\cdot | t_i)$. The set of finite common-prior types is not dense in the set of finite types, and thus is not dense in the set of all types. Intuitively, this is because the strategic implications of a common prior (such as certain no-trade theorems) do not extend to general types. Since Lipman (2003) shows that the set of finite common prior types is dense in the product topology,

this observation demonstrates a general distinction between the product topology and the strategic topology.

To demonstrate this observation, it is enough to give a single game that separates common prior types from non common-prior types. Suppose there are two states, L and R . Each player has two actions, Y (trade) and N (no trade). If a player says N she gets zero, if they both say Y they get $(1, -2)$ and $(-2, 1)$ in states L and R respectively (a negative-sum trade) and in case one says Y and the other says N the one who said Y gets -2 (independent of the state). Clearly both players saying Y is 0-rationalizable for some non-common prior type: if a player believes there is common certainty that each player believes the state is the one favorable to him, they may both say Y . But with common priors, the only ε -rationalizable action, for ε small enough, is N : Let \tilde{T} be the set of common-prior type pairs for whom Y is $\frac{1}{4}$ -rationalizable; for each $t_1 \in \tilde{T}_1$, the probability of state L and \tilde{T} has to be at least $\frac{7}{12}$ (since $\frac{7}{12} - 2(\frac{5}{12}) = -\frac{1}{4}$). However, a similar property, replacing R for L , must also hold for all $t_2 \in \tilde{T}_2$, so there cannot be a common prior on $\tilde{T} \times \{L, R\}$ with such conditionals.

7.4 *Alternative solution concepts*

We used the solution concept of interim correlated rationalizability to define the strategic topology. We noted two reasons for doing this: the set of interim correlated rationalizable actions depend only on hierarchies of beliefs and the solution concept captures the implications of common knowledge of rationality.

One might wonder what would happen with alternative solution concepts, such as Nash equilibrium or interim independent rationalizability. But the set of Nash equilibrium actions or interim-independent-rationalizable actions depend in general not just on the belief hierarchies but also on “redundant” types: those that differ in their ability to correlate their behavior with others’ actions and states of the world. In defining a topology for these solution concepts, one would have to decide what to do about this dependence.

Suppose one wanted to define a topology on hierarchies of beliefs (despite the fact that hierarchies of beliefs do not determine these other solution concepts). One approach would be to examine all actions that could be played by any type with a given hierarchy of beliefs under that solution concept (allowing for all possible type spaces and not only the universal type space as we do here). Dekel et al. (2006) show that, given any game and fixed hierarchy of beliefs, the union—over all type spaces that contain a type with that hierarchy of beliefs—of the equilibrium actions equals the interim-correlated-rationalizable actions for that hierarchy of beliefs. The same arguments can be used to prove the same conclusion for interim-independent-rationalizability.

Another approach would be to fix a solution concept and construct a larger representation of types that included hierarchies of beliefs but also incorporated the redundant types that are relevant for the solution concept. One would then construct a topology on this larger space. The first part of this approach—constructing the larger type space that incorporates the redundant types relevant for the solution concept—is carried out by Ely and Pęski (2006) for the solution concept of interim-independent-rationalizability for two-person games.

7.5 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. As we said earlier, when the number of actions is unbounded, there is usually a metric on actions and accompanying constraints on the set of admissible pay-off functions. Uniformity over all games without such restrictions seems too demanding and hence of less interest. Despite this, it seems useful to understand how our results would change if we did ask for uniformity over games.

Let $A_i(G)$ denote the set of actions for player i in a given game G . A distance on types that is uniform in games is:

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

This metric yields a topology that is finer than that induced by the metric d , so the topology is finer than necessary for the upper and lower convergence properties that we took as our goal. Finite types are not dense with this topology. To show this we use the fact that convergence in this topology is equivalent to convergence in the following uniform topology on beliefs:

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

where F_k is the collection of bounded functions mapping $T^* \times \Theta$ that are measurable with respect to k^{th} level beliefs.

PROPOSITION 2. *The metrics d^* and d^{**} are equivalent.*

An argument of [Morris \(2002\)](#) implies that finite types are not dense in the uniform topology on beliefs.¹⁹ Together with the preceding proposition this implies that finite types are not dense in the uniform strategic topology generated by d^* .

7.6 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 ([Cr mer 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.²⁰

¹⁹For a fixed random variable on payoff states, we can identify the higher-order expectations of a type, i.e., the expectation of the random variable, the expectation of the other player's expectation of the random variable, and so on. Convergence in the metric d^{**} implies that there is uniform convergence of these higher order expectations. [Morris \(2002\)](#) shows that finite types are not dense in a topology defined in terms of uniform convergence of higher order expectations. Thus finite types are not dense under d^{**} .

²⁰A result in [Bergemann and Morris \(2001\)](#) (which does not appear in the published version of the paper, [Bergemann and Morris 2005](#)), shows that both the set of full-surplus-extraction types and the set of non-

7.7 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 do 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 (2006) on the drawbacks of analyzing genericity 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 Crémer and McLean (1985), McAfee and Reny (1992), and Jehiel and Moldovanu (2001).

APPENDIX

For some proofs we use an alternative characterization of the interim-correlated-rationalizable actions.

DEFINITION 10. Fix a game $G = (A_i, g_i)_{i \in I}$ and a belief-closed subspace $(T_i, \pi_i^*)_{i \in I}$. Given $S = (S_1, S_2)$, where each $S_i : T_i \rightarrow 2^{A_i} \setminus \emptyset$, we say that S is an ε -best response set if $S_i(t_i) \subseteq BR_i(t_i, S_j, G, \varepsilon)$.

In Dekel et al. (2006) we prove that if S is a 0-best response set then $S_i(t_i) \subseteq R(t_i, G, 0)$; the extension to positive ε is immediate.

Let

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

LEMMA 9. For any integer m and $\varepsilon > 0$, there exists a finite collection of m action games $\overline{\mathcal{G}}^m$ such that, for every $g \in \mathcal{G}^m$, there exists $g' \in \overline{\mathcal{G}}^m$ such that $D(g, g') \leq \varepsilon$.

PROOF. Assume without loss of generality that M is an integer. 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 $D(g, g') \leq 1/(2N)$. \square

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.

LEMMA 10. For all $i, t_i, m, a_i \in A_i^m$, and $g, g' \in \mathcal{G}^m$

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

PROOF. 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 $h_i(t_i | a_i, g') + 2D(g, g') \geq h_i(t_i | a_i, g)$. \square

LEMMA 11. (i) If for each m and each $g \in \mathcal{G}^m$ there exists $\hat{\varepsilon}^n \rightarrow 0$ such that $h_i(t_i | a_i, g) < h_i(t_i^n | a_i, g) + \hat{\varepsilon}^n$ for every n and $a_i \in A_i^m$, then for each m there exists $\varepsilon^n \rightarrow 0$ such that $h_i(t_i | a_i, g) < h_i(t_i^n | a_i, g) + \varepsilon^n$ for every $n, g \in \mathcal{G}^m$, and $a_i \in A_i^m$.

(ii) If for each m and each $g \in \mathcal{G}^m$ there exists $\hat{\varepsilon}^n \rightarrow 0$ such that $h_i(t_i^n | a_i, g) < h_i(t_i | a_i, g) + \hat{\varepsilon}^n$ for every n and $a_i \in A_i^m$, then for each m there exists $\varepsilon^n \rightarrow 0$ such that $h_i(t_i^n | a_i, g) < h_i(t_i | a_i, g) + \varepsilon^n$ for every $n, g \in \mathcal{G}^m$, and $a_i \in A_i^m$.

PROOF. Lemma 10 implies that for fixed m, h_i is continuous in g . Assume now to the contrary that part (i) was false. Then there exists m and $\delta > 0$ such that for all n there is $g^n \in \mathcal{G}^m$ with $h_i(t_i | a_i, g^n) \geq h_i(t_i^n | a_i, g^n) + \delta$. Since \mathcal{G}^m is a compact metric space, the sequence g^n has a convergent sub-sequence; denote the limit of that subsequence by \bar{g} . Then $h_i(t_i | a_i, \bar{g}) \geq h_i(t_i^n | a_i, \bar{g}) + \delta/2$, contradicting the hypothesis. The same argument proves part (ii). \square

LEMMA 7. For any integer m and $\varepsilon > 0$, there exists a finite collection of m action games $\overline{\mathcal{G}}^m$ such that, for every $g \in \mathcal{G}^m$, there exists $g' \in \overline{\mathcal{G}}^m$ such that for all i, t_i , and $a_i \in A_i^m$

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

PROOF. By Lemma 9, we can choose a finite collection of games $\overline{\mathcal{G}}^m$ such that, for every $g \in \mathcal{G}^m$, there exists $g' \in \overline{\mathcal{G}}^m$ such that $D(g, g') \leq \varepsilon/2$. Lemma 10 now implies that we also have

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

and

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

LEMMA 12. Fix any finite collection of m action games $\overline{\mathcal{G}}^m$ and $\delta > 0$. There exists a finite type space $(\tilde{T}_i, \tilde{\pi}_i)_{i=1,2}$ and functions $(f_i)_{i=1,2}$, each $f_i : T^* \rightarrow \tilde{T}_i$, 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. Let $(\tilde{T}_i, \tilde{\pi}_i)$ and $(f_i)_{i=1,2}$ be as constructed in Section 6.2. Fix $\varepsilon \in \{0, \delta, \dots, (2M)^\delta\}$. Let $S_i(\tilde{t}_i) = R_i(\zeta_i(\tilde{t}_i), g, \varepsilon)$.

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

$$a_i \in R_i(t_i, g, \varepsilon)$$

and let $\tilde{t}_i = f_i(t_i)$. Because $\varepsilon \in \{0, \delta, \dots, (2M)^\delta\}$, we have $R_i(t_i, g, \varepsilon) = R_i(\zeta_i(\tilde{t}_i), g, \varepsilon)$ and thus

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

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

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

Now define $\tilde{\nu} \in \Delta(\tilde{T}_j \times \Theta \times A_j)$ by

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

By construction,

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

So $a_i \in BR_i(S)(\tilde{t}_i)$.

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

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

LEMMA 13. Fix any finite collection of m action games $\overline{\mathcal{G}}^m$ and $\delta > 0$. There exists a finite type space $(\tilde{T}_i, \tilde{\pi}_i)_{i=1,2}$ 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 $t_i, g \in \overline{\mathcal{G}}^m$, and $a_i \in A_i^m$.

PROOF. We use the type space from [Lemma 12](#), which has the property that

$$R_i(t_i, g, \varepsilon) \subseteq R_i(f_i(t_i), g, \varepsilon) \tag{3}$$

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,

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

By (3),

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

Thus

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

LEMMA 8. 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 $f_i : T^* \rightarrow \tilde{T}_i$, 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 $a_i \in A_i^m$.

PROOF. Fix m and $\xi > 0$. By Lemma 7, there exists a finite collection of m -action games $\overline{\mathcal{G}}^m$ such that for every finite-action game g , there exists $g' \in \overline{\mathcal{G}}^m$ such that

$$h_i(t_i | a_i, g) \leq h_i(t_i | a_i, g') + \frac{1}{3}\xi \quad \text{and} \quad h_i(t_i | a_i, g') \leq h_i(t_i | a_i, g) + \frac{1}{3}\xi \quad (4)$$

for all i, t_i , and a_i . By Lemma 13, there exists a finite type space $(\tilde{T}_i, \tilde{\pi}_i)_{i=1,2}$ and functions $(f_i)_{i=1,2}$, each $f_i : T^* \rightarrow \tilde{T}_i$, such that

$$h_i(f_i(t_i) | a_i, g) \leq h_i(t_i | a_i, g) + \frac{1}{3}\xi \quad (5)$$

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

Now fix any i, t_i, a_i and g . By (4), there exists g' such that

$$h_i(t_i | a_i, g') - h_i(t_i | a_i, g) \leq \frac{1}{3}\xi$$

and

$$h_i(f_i(t_i) | a_i, g) - h_i(f_i(t_i) | a_i, g') \leq \frac{1}{3}\xi.$$

By (5),

$$h_i(f_i(t_i) | a_i, g') - h_i(t_i | a_i, g') \leq \frac{1}{3}\xi.$$

So

$$\begin{aligned} & h_i(f_i(t_i) | a_i, g) - h_i(t_i | a_i, g) \\ & \leq (h_i(f_i(t_i) | a_i, g) - h_i(f_i(t_i) | a_i, g')) + (h_i(f_i(t_i) | a_i, g') - h_i(t_i | a_i, g')) \\ & \quad + (h_i(t_i | a_i, g') - h_i(t_i | a_i, g)) \\ & \leq \xi. \end{aligned} \quad \square$$

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 that 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 $\pi_i : \{1, \dots, n\} \rightarrow \Delta(\{1, \dots, n\} \times \Theta)$ (as before these are to be viewed as a belief-closed subspace of the universal type space.) 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 requires 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).

We now define a sequence of type spaces $((T_i^k, \pi_i^k)_{i \in I})$ for $k = 1, 2, \dots, \infty$. Let us suppose each player i observes an additional signal $z_i \in \{1, 2, \dots\}$, and define $T_i^k =$

$T_i \times \{1, 2, \dots\}$, with typical element (n_i, z_i) .) Fix $\lambda \in (0, 1)$ and for each $k = 1, 2, \dots$, choose $\pi_i^k : T_i^k \rightarrow \Delta(T_j^k \times \Theta)$ to satisfy the following two properties:

$$\left| \pi_i^k((n_j, z_j), \theta \mid (n_i, z_i)) - (1 - \lambda)\lambda^{z_j-1} \pi_i(n_j, \theta \mid n_i) \right| \leq \frac{1}{k} \tag{6}$$

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

$$\pi_i^k(\theta \mid (n_i, z_i)) \neq \pi_i^k(\theta' \mid (n'_i, z'_i)) \tag{7}$$

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

For $k = \infty$, set

$$\pi_i^\infty((n_j, z_j), \theta \mid (n_i, z_i)) = (1 - \lambda)\lambda^{z_j-1} \pi_i(n_j, \theta \mid n_i),$$

where (7) is not satisfied (but holds instead with equality).

We distinguish the different copies of T_i^k with the superscript k because we identify a type (n_i, z_i) in T_i^k as potentially distinct from $(n_i, z_i) \in T_i^{k'}$ when viewed as types in the universal type space. Note that in the type space $(T_i^\infty, \pi_i^\infty)$, for each $n_i \in T_i$ every type $(n_i, z_i) \in T_i \times \{1, 2, \dots\}$ corresponds to the same type in the universal type space, namely the type in the universal type space that corresponds to type n_i in the type space (T_i, π_i) . On the other hand, from (7) we see that for all other type spaces with $k \neq \infty$, each distinct pair of types $(n_i, z_i) \neq (n'_i, z'_i)$ in (T_i^k, π_i^k) corresponds to distinct types in the universal type space. Let $t^k \in T^*$ be the type in the universal type corresponding to type $(1, 1)$ in the type space $(T_i^k, \pi_i^k)_{i=1,2}$.

Now (7) also implies that each $t^k \notin T_F$.

We argue that the sequence t^k converges to t^* in the strategic topology. To see why, for any $(n_i, z_i) \in T_i^k = T_i \times \{1, 2, \dots\}$, 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 $(T_i^\infty, \pi_i^\infty)_{i=1,2}$. This is true because, as noted, the type space $(T_i^\infty, \pi_i^\infty)_{i=1,2}$ and the original type space $(T_i, \pi_i)_{i=1,2}$ correspond to the same belief-closed subspace of the universal type space. (The only difference is that in $(T_i^\infty, \pi_i^\infty)_{i=1,2}$ it is common knowledge that each player i observes a conditionally independent draw with probabilities $(1 - \lambda)\lambda^{z_i-1}$ on $\{1, 2, \dots\}$.) But now by (6), S is an $\eta + 2M/k$ best response set for G . Thus if $a_i \in R_i(t^*, G, \eta)$, then $a_i \in R_i(t^k, G, \eta + 2M/k)$. Thus the sequence (t^k, t^*) satisfies the lower strategic convergence property. By **Corollary 1**, this implies strategic convergence. By **Theorem 1**, we also have product convergence. \square

PROPOSITION 2. *d^{**} is equivalent to d^* .*

PROOF. First observe that if $d^{**}(t_i, t'_i) \leq \epsilon$, then for any measurable $f : T_j^* \times \Theta \rightarrow [-M, M]$,

$$\left| E(f \mid \pi^*(t_i)) - E(f \mid \pi^*(t'_i)) \right| \leq 2\epsilon. \tag{8}$$

To see this, for any k , we define $f_k : T_j^* \times \Theta \rightarrow [-M, M]$ that is measurable with respect to i 's k^{th} level beliefs. Let

$$f_k(\bar{t}_j, \theta) = \int_{\{(t_j, \theta) : (\hat{\pi}_{j1}(\bar{t}_j))_{l=0}^{k-1} = (\hat{\pi}_{j1}(t_j))_{l=0}^{k-1}\}} f(t_j, \theta) d(\pi^*(t_i)).$$

That is, f_k is the expected value according to $\pi^*(t_i)$ of the function f evaluated over all t_j with the same first k levels. Now $f_k \rightarrow f$ pointwise so by the bounded convergence theorem, $E(f_k | \pi^*(t'_i)) \rightarrow E(f | \pi^*(t'_i))$. By the definition (and iterated expectations) $E(f_k | \pi^*(t_i)) = E(f | \pi^*(t_i))$. Since $d^{**}(t_i, t'_i) \leq \varepsilon$ we know that for all k , $|E(f_k | \pi^*(t_i)) - E(f_k | \pi^*(t'_i))| \leq \varepsilon$.

Now suppose that $a_i \in R_i(t_i, G, \delta)$. Then there exists $\nu \in \Delta(T_j^* \times \Theta \times A_j)$ such that

$$\begin{aligned} \nu \left[\{(t_j, \theta, a_j) : a_j \in R_j(t_j, G, \delta)\} \right] &= 1 \\ \text{marg}_{T_j^* \times \Theta} \nu &= \pi_i^*(t_i) \\ \int_{(t_j, \theta, a_j)} [g_i(a_i, a_j, \theta) - g_i(a'_i, a_j, \theta)] d\nu &\geq -\delta \text{ for all } a'_i \in A_i. \end{aligned}$$

Let ν' be a measure whose marginal on $T_j^* \times \Theta$ is $\pi_i^*(t'_i)$ and whose probability on A_j , conditional on (t_j, θ) , is the same as ν . Since $T_j^* \times \Theta \times A_j$ is a separable standard measure space there exist conditional probabilities (see Parthasarathy 1967, Theorem 8.1) $\nu(\cdot | T_j^* \times \Theta) \in \Delta(A_{-i})$, measurable as a function of $T_j^* \times \Theta$. Define $\nu' \in \Delta(T_j^* \times \Theta \times A_{-i})$, by setting, for measurable $F \subset T_j^* \times \Theta$, $\nu'(F \times \{a_j\}) = \int_F (\nu(a_j | t_j, \theta) \times \pi_i(t_i)) [dt_j, \theta]$.²¹ Let f_{a_i, a'_i} be a function taking the value $\int_{A_j} (g_i(a_i, a_j, \theta) - g_i(a'_i, a_j, \theta)) d\nu(a_j | T_j^* \times \Theta)$ at each (t_j, θ) . So by (8),

$$\begin{aligned} \int_{(t_j, \theta, a_j)} [g_i(a_i, a_j, \theta) - g_i(a'_i, a_j, \theta)] d\nu' \\ \geq \int_{(t_j, \theta, a_j)} [g_i(a_i, a_j, \theta) - g_i(a'_i, a_j, \theta)] d\nu - 2\varepsilon(2) \geq -\delta - 4\varepsilon. \end{aligned}$$

Thus $a_i \in R_i(t'_i, G, \delta + 4\varepsilon)$. Since this argument holds for every game G independent of the cardinality of the action sets, we have $d^*(t_i, t'_i) \leq 4\varepsilon$.

On the other hand, if $d^{**}(t_i, t'_i) \geq \varepsilon$ then there exists k such that

$$|E(f | \pi^*(t_i)) - E(f | \pi^*(t'_i))| \geq \frac{1}{2}\varepsilon$$

for some bounded f that is measurable with respect to k^{th} level beliefs. Now Lemma 4 states that that we can construct a game G and action a_i such that $h_i(t'_i | a_i, G) + \frac{1}{2}\varepsilon < h_i(t_i | a_i, G)$. Thus $d^*(t_i, t'_i) \geq \frac{1}{2}\varepsilon$.

²¹A similar construction appears in Dekel et al. (2006).

We conclude that

$$\frac{1}{2}d^{**}(t_i, t'_i) \leq d^*(t_i, t'_i) \leq 4d^{**}(t_i, t'_i)$$

or equivalently

$$\frac{1}{4}d^*(t_i, t'_i) \leq d^{**}(t_i, t'_i) \leq 2d^*(t_i, t'_i).$$

Thus $d^{**}(t_i^n, t_i) \rightarrow 0$ if and only if $d^*(t_i^n, t_i) \rightarrow 0$. □

REFERENCES

- Bergemann, Dirk and Stephen Morris (2001), “Robust mechanism design.” Technical report, Princeton University. [301]
- Bergemann, Dirk and Stephen Morris (2005), “Robust mechanism design.” *Econometrica*, 73, 1771–1813. [301]
- Brandenburger, Adam and Eddie Dekel (1987), “Rationalizability and correlated equilibria.” *Econometrica*, 55, 1391–1402. [276]
- Brandenburger, Adam and Eddie Dekel (1993), “Hierarchies of beliefs and common knowledge.” *Journal of Economic Theory*, 59, 189–198. [281, 286]
- Cr mer, Jacques and Richard McLean (1985), “Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent.” *Econometrica*, 53, 345–361. [301, 302]
- Dekel, Eddie, Drew Fudenberg, and David K. Levine (2004), “Learning to play Bayesian games.” *Games and Economic Behavior*, 46, 282–303. [278]
- Dekel, Eddie, Drew Fudenberg, and Stephen Morris (2006), “Interim rationalizability.” Unpublished paper, Harvard University. [276, 278, 281, 283, 284, 285, 300, 302, 307]
- Ely, Jeffrey C. and Marcin Pełski (2006), “Hierarchies of belief and interim rationalizability.” *Theoretical Economics*, 1, 19–65. [296, 300]
- Fudenberg, Drew and David K. Levine (1986), “Limit games and limit equilibria.” *Journal of Economic Theory*, 38, 261–279. [277]
- Geanakoplos, John D. and Heraklis M. Polemarchakis (1982), “We can’t disagree forever.” *Journal of Economic Theory*, 28, 192–200. [276]
- Harsanyi, John C. (1967–68), “Games with incomplete information played by ‘Bayesian’ players, I–III.” *Management Science*, 14, 159–182, 320–334, 486–502. [275]
- Heifetz, Aviad (1993), “The Bayesian formulation of incomplete information—the non-compact case.” *International Journal of Game Theory*, 21, 329–338. [281]
- Heifetz, Aviad and Zvika Neeman (2006), “On the generic (im)possibility of full surplus extraction in mechanism design.” *Econometrica*, 74, 213–233. [299, 302]

Heifetz, Aviad and Dov Samet (1998), "Topology-free typology of beliefs." *Journal of Economic Theory*, 82, 324–341. [281, 282]

Jehiel, Philippe and Benny Moldovanu (2001), "Efficient design with interdependent valuations." *Econometrica*, 69, 1237–1259. [302]

Kajii, Atsushi and Stephen Morris (1997), "Payoff continuity in incomplete information games." *Journal of Economic Theory*, 82, 267–276. [278]

Lipman, Barton L. (2003), "Finite order implications of common priors." *Econometrica*, 71, 1255–1267. [277, 286, 299]

McAfee, R. Preston and Philip R. Reny (1992), "Correlated information and mechanism design." *Econometrica*, 60, 395–421. [301, 302]

Mertens, Jean-François, Sylvain Sorin, and Shmuel Zamir (1994), "Repeated games: Part a. background material." Discussion Paper 9420, CORE, Louvain-la-Neuve. [276, 281]

Mertens, Jean-François and Shmuel Zamir (1985), "Formulation of Bayesian analysis for games with incomplete information." *International Journal of Game Theory*, 14, 1–29. [275, 278, 281, 282, 286, 298]

Monderer, Dov and Dov Samet (1996), "Proximity of information in games with common beliefs." *Mathematics of Operations Research*, 21, 707–725. [278, 296]

Morris, Stephen (2002), "Typical types." Unpublished paper, Princeton University. [301]

Munkres, James R. (1975), *Topology*. Prentice-Hall, Englewood Cliffs, New Jersey. [288]

Neeman, Zvika (2004), "The relevance of private information in mechanism design." *Journal of Economic Theory*, 117, 55–77. [302]

Parthasarathy, K. S. (1967), *Probability Measures on Metric Spaces*. Academic Press, New York. [307]

Rubinstein, Ariel (1989), "The electronic mail game: Strategic behavior under 'almost common knowledge'." *American Economic Review*, 79, 385–391. [276, 279]

Weinstein, Jonathan and Muhamet Yildiz (2003), "Finite order implications of any equilibrium." Unpublished paper, Department of Economics, MIT. [277, 280, 286]

Yildiz, Muhamet (2006), "Generic uniqueness of rationalizable outcomes." Unpublished paper, Department of Economics, MIT. [277, 280]