Robust Multiplicity with a Grain of Naiveté[*]

Aviad Heifetz[†]   Willemien Kets[‡]

December 11, 2013

Abstract

In an important paper, Weinstein and Yildiz (2007) show that if players have an infinite depth of reasoning and this is commonly believed, types generically have a unique rationalizable action in games that satisfy a richness condition. We show that this result does not extend to environments where players may have a finite depth of reasoning, or think it is possible that the other player has a finite depth of reasoning, or think that the other player may think that is possible, and so on, even if this so-called "grain of naiveté" is arbitrarily small. More precisely, we show that even if there is almost common belief in the event that players have an infinite depth of reasoning, there are types with multiple rationalizable actions, and the same is true for "nearby" types. Our results demonstrate that both uniqueness and multiplicity are robust phenomena when we relax the assumption that it is common belief that players have an infinite depth, if only slightly.

[†]Department of Management and Economics, Open University of Israel. E-mail: aviadhe@openu.ac.il. Phone: +972-9-778-1878.

[‡]Kellogg School of Management, Northwestern University. E-mail: w-kets@kellogg.northwestern.edu. Phone: +1-505-204 8012.

1 Introduction

Multiplicity of predictions naturally arises in many situations of economic interest, ranging from bank runs, currency attacks, debt crises, to arms races. One could view models with multiple predictions as formalizing the idea that strategic uncertainty plays a critical role in such games, so that there is an inherent indeterminacy of outcomes.[1] An alternative perspective is to view models with multiple predictions as incomplete theories, where the multiplicity can be resolved by extending the model; for example, by introducing incomplete information. A striking result, due to Carlsson and van Damme (1993), is that a coordination game has an essentially unique prediction when players observe the state of nature with some small noise; a type has multiple rationalizable actions only if its signal is on the boundary of a set of signals for which the rationalizable action is unique.[2] Weinstein and Yildiz (2007) formalize and strengthen this idea: in any game, and for any information structure, types generically have a unique rationalizable outcome, and multiplicity occurs if and only if we are in a knife-edge case.[3]

An assumption that plays an essential role in the proof of these results is that players have an infinite depth of reasoning: not only are players able to think about the payoffs of the game and form beliefs about these payoffs, they can also form beliefs about others' beliefs about payoffs, about the others' beliefs about *their* opponents' beliefs, and so on, *ad infinitum.* In fact, not only are players assumed to have an infinite depth of reasoning, it is assumed that this is common belief. This is, of course, at best an idealization: some agents may only have a finite depth of reasoning; others may have an infinite depth of reasoning, but may think that their opponent has a finite depth, and so on.[4]

We therefore relax the assumption that it is common belief that players have an infinite depth of reasoning, and show that the results of Carlsson and van Damme (1993) and Weinstein and Yildiz (2007) do not extend to this environment. More specifically, we show that multiplicity of rationalizable outcomes is a typical and robust phenomenon if we do not insist

---

[1]See, e.g., Diamond and Dybvig (1983), Obstfeld (1986, 1996), and Cole and Kehoe (2000).

[2]This result has inspired a large literature with many important applications. See, e.g., Morris and Shin (1998, 2004), Baliga and Sjöström (2004), Corsetti et al. (2004), and Goldstein and Pauzner (2005). See Morris and Shin (2003) for a review.

[3]The results of Carlsson and van Damme and Weinstein and Yildiz presume that there is no a priori restriction on payoffs; see Section 5.

[4] An important literature in experimental literature suggests that subjects indeed only use a finite depth of reasoning in a range of games. See, e.g., Nagel (1995), Stahl and Wilson (1995), Ho et al. (1998), Costa-Gomes et al. (2001), Camerer et al. (2004), and Crawford and Iriberri (2007); see Heinemann et al. (2004, 2009), Cabrales et al. (2007) for experimental studies of coordination games with noisy information, with a focus on the effect of subjects' depth of reasoning. See Crawford et al. (2012) for a survey.

on players having common belief in an infinite depth. This is true even if players have an infinite depth of reasoning and it is almost common belief that players have an infinite depth.

To establish our results, we need a framework to model the beliefs of players with finite and infinite depth. Following Mertens and Zamir (1985), we construct a space of belief hierarchies, and show that this space can be used to define a so-called universal type space that contains all beliefs in a sense we make precise. We show that the universal type space of Mertens and Zamir forms a (strict) subset of our universal space that is characterized by the event that players have an infinite depth of reasoning and have common belief in that event (Proposition 3.4). Our universal type space thus provides the appropriate environment to study the effects of relaxing the assumption that players have an infinite depth of reasoning and have common belief in infinite depth.

Our first main result (Theorem 5.2) shows that for generic games in the class of games studied by Carlsson and van Damme, there are types with an infinite depth of reasoning and with almost common belief in infinite depth that have multiple rationalizable actions; moreover, "nearby" types also have multiple rationalizable actions. Thus, there is robust multiplicity in generic games even if there is almost common belief that players have an infinite depth.

To see the intuition, it is useful to consider a simple example. Consider the game in Figure 1, taken from Morris and Shin (2003). When payoffs are commonly known, the game can have multiple rationalizable outcomes: if it is commonly known that $s = \frac{1}{2}$, for example, both actions are rationalizable for each player. However, as shown by Carlsson and van Damme, if beliefs are perturbed slightly, so that a player believes that $s = \frac{1}{2}$, believes that the other believes that, and so on, for a large but finite number n of iterations, but believes that the other believes... ($n + 1$ times)... that the other believes that $s = -1$, then her unique rationalizable action is b_i.

	a_2	b_2
a_1	s, s	$s - 1, 0$
b_1	$0, s - 1$	$0, 0$

Figure 1: A game in which payoffs depend on the state of nature s.

Now consider a player, say player 1, who forms a belief about payoffs, but not about player 2's beliefs about payoffs; in that case, we say that player 1 has depth of reasoning equal to 1. A plausible assumption is that a player with depth 1 does not rule out any action of her opponent.[5] Suppose player 1 assigns equal probability to $s = 2$ and $s = -1$. Then, under

---

[5]Indeed, it is common in the experimental literature (footnote 4) to assume that a player with depth 1

3

the conjecture that her opponent plays a_2, it is optimal for her to choose a_1. On the other hand, under the conjecture that player 2 chooses b_2, it is a best response for her to choose b_1. It follows that both actions are rationalizable for player 1 under such a belief. We show that this remains true for types with slightly perturbed beliefs.

Next suppose a player has depth 2 and thinks that his opponent has depth 1. If he thinks that his opponent assigns roughly equal probability to $s = 2$ and $s = -1$, then, by the above argument, he cannot rule out any action of his opponent. If this player in turn assigns equal probability to $s = 2$ and $s = -1$, then, by a similar argument as above, multiple actions are rationalizable for that player, as well as for nearby types.

Theorem 5.2 demonstrates that for generic games, this simple argument can be extended to more complex environments to show that even types with an infinite depth and with almost common belief in infinite depth, have multiple rationalizable actions, and the same is true for types with similar beliefs. This result implies that the result of Weinstein and Yildiz that types generically have a unique rationalizable action when there is common belief in infinite depth does not carry over to this more general environment, where players can have a finite depth of reasoning, or may think that others have a finite depth of reasoning, and so on (Corollary 5.3).

Our second main result focuses specifically on one of the canonical global games discussed in the literature (Carlsson and van Damme, 1993; Morris and Shin, 2003). In such games, players receive a noisy private signal about the state of nature. While Theorem 5.2 implies that there are types that have multiple rationalizable actions in this game, the beliefs of the types constructed in the proof are generally not consistent with this information structure with noisy signals. Theorem 5.4 shows that in fact, there is robust multiplicity even if beliefs are required to be consistent with this information structure, players have an infinite depth and almost common belief in the infinite depth, and have signals that almost makes one of the actions (uniquely) dominant, provided that the noise level is sufficiently small. As we discuss in Section 5, the difficulty in proving this result is to show that there is a positive noise level for which a type has multiple rationalizable actions, even for extreme signals and very high levels of mutual belief in infinite depth.

For this particular game, Theorem 5.4 is stronger than Theorem 5.2. Theorem 5.4 says that we can always find a type whose belief is consistent with the information structure with multiple rationalizable actions (and that has an infinite depth and almost common belief in an infinite depth), while the beliefs of the types in Theorem 5.2 can be arbitrary. On the other hand, by fixing the information structure, Theorem 5.4, by its nature, can apply only

thinks that her opponent chooses an action uniformly at random, unless some action is particularly salient (Crawford et al., 2012).

4

to a particular game, whereas Theorem 5.2 applies to generic games in a broad class. Thus, the two results complement each other.

The intuition behind our results is that players with a finite depth of reasoning provide a "seed" of multiplicity, as in the example above. By contrast, it follows from the results from Weinstein and Yildiz that no such seed exists when there is common belief in the event that players have an infinite depth of reasoning. For example, if the set of states of nature is $S = \{-1, 2\}$, then, under the assumption that there is common belief in an infinite depth of reasoning, a player has multiple rationalizable actions if and only if it is (correct) common belief that players assign a probability sufficiently close to $\frac{1}{2}$ to $s = 2$. However, this is a knife-edge case.

We show that this seed of multiplicity can be used to construct an open set of types with multiple rationalizable actions even under strong restrictions on players' beliefs, such as almost common belief in the event that players have an infinite depth of reasoning, or that beliefs be consistent with a particular information structure. We thus use a similar proof technique as the existing literature that obtains unique predictions by employing the "existence...of dominance solvable games that serve as take-offs for the...argument, and, thus, exert a kind of remote influence" (Carlsson and van Damme, 1993, p. 992). Rather than dominance solvable games, we use types with a finite depth of reasoning as a starting point, and rather than showing generic uniqueness, we demonstrate that multiplicity can be robust.

The remainder of this paper is organized as follows. After discussing some preliminaries, Section 3 defines belief hierarchies of finite and infinite depth and constructs the universal type space. Section 4 defines games with incomplete information, and Section 5 presents our main results. Section 6 discusses the related literature, and the online appendix contains additional results.

2   Preliminaries

We follow the standard conventions for subspaces, products, and (disjoint) unions of topological spaces. A subspace of a topological space is endowed with the relative topology, and the product of a collection of topological spaces is endowed with the product topology. If $(V_\lambda)_{\lambda \in \Lambda}$ is a family of topological spaces (possibly made disjoint by replacing some V_λ with a homeomorphic copy), then $\bigcup_\lambda V_\lambda$ is endowed with the sum topology, that is, a subset $U \subseteq \bigcup_{\lambda \in \Lambda} V_\lambda$ is open in $\bigcup_{\lambda \in \Lambda} V_\lambda$ if and only if $U \cap V_\lambda$ is open in V_λ for each $\lambda \in \Lambda$. In particular, if V_λ, $\lambda \in \Lambda$, is a metric space with metric d_λ (bounded by 1), then the sum topology on $\bigcup_{\lambda \in \Lambda} V_\lambda$ is
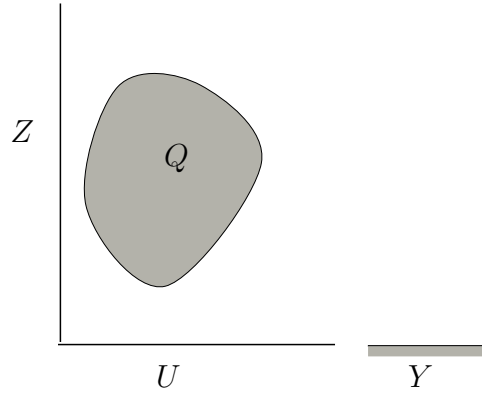
Figure 2: The space W (shaded gray) is the union of $Q \subseteq U \times Z$ and of Y. The space V is the union of U and Y.

metrized by the metric d_Λ, defined by

$$d_\Lambda(v, v') := \begin{cases} d_\lambda(v, v') & \text{if } v, v' \in V_\lambda; \\ 1 & \text{otherwise.} \end{cases}$$

Hence, if v, v' belong to different component spaces $V_\lambda, V_{\lambda'}$, then they are not close in the sum topology. The sum of a countable collection of Polish spaces is Polish (Kechris, 1995, Prop. 3.3).

Given a topological space V, denote by $\Delta(V)$ the set of probability measures on the Borel σ-algebra $\mathscr{B}(V)$ on V. We endow $\Delta(V)$ with the topology of weak convergence, and with its associated Borel σ-algebra. If V is Polish, then this σ-algebra is generated by the sets

$$\{\mu \in \Delta(V) : \mu(E) \geq p\} \qquad E \in \mathscr{B}(V), p \in [0, 1];$$

a result that we will use often without mention.

We extend the definition of a marginal probability measure to a union of measurable spaces. Let V be the union of the disjoint sets U and Y, and let $Q \subseteq U \times Z$ and $W = Q \cup Y$, where all spaces are assumed to be topological spaces; see Figure 2. Then for $\mu \in \Delta(W)$ denote by $\text{marg}_V \mu \in \Delta(V)$ the probability measure defined by

$$\text{marg}_V \mu(E) = \mu(\{(u, z) \in Q : u \in E\}) + \mu(E \cap Y)$$

for every measurable set $E \subseteq V$. If μ is a probability measure on a product space $U \times Y$, and E is a measurable subset of U, then we sometimes write $\mu(E)$ for $\text{marg}_U \mu(E)$.

3   Belief hierarchies

This section constructs the space of all belief hierarchies of finite and infinite depth. We then define the universal type space for players with a finite or infinite depth of reasoning, analogously to the universal Harsanyi type space of Mertens and Zamir (1985). The proofs for the results in this section can be found in the online appendix.

6

3.1 Construction

This section constructs the space of all coherent belief hierarchies. There is a fixed, finite set of players N. There is a set S of *states of nature*. Each player $i \in N$ can have private information about the state of nature: she may receive a *signal* $x_i \in X_i$. A belief hierarchy for player i specifies her belief about the state of nature and the other players' signals, i.e., about $S \times X_{-i}$, her beliefs about her opponents' beliefs, and so on, up to some finite or infinite order.[6] That is, each hierarchy is associated with a *depth* of reasoning; a hierarchy's depth can either be finite or infinite.[7]

Formally, assume that S and X_i, $i \in N$, are compact metric.[8] We construct two sequences of spaces for each player i, H_i^m and $\widetilde{H}_i^m$, $m \geq 0$, with H_i^m the set of *mth-order belief hierarchies* that "stop" reasoning at order m, and $\widetilde{H}_i^m$ the set of belief hierarchies that "continue" to reason at that order. The belief hierarchies in H_i^m form the belief hierarchies with depth of reasoning equal to m, while the belief hierarchies in $\widetilde{H}_i^m$ are used to construct the belief hierarchies of depth at least $m+1$ (and possibly infinite).

For a player $i \in N$, fix two arbitrary labels $h_i^{*,0}$ and μ_i^0; the label $h_i^{*,0}$ is used to define the set H_i^0 of belief hierarchies of players who do not reason at all, while $\tilde{\mu}_i^0$ is just a notational "seed" on which the hierarchies of more sophisticated types will be built. Let $\widetilde{H}_i^0 := X_i \times \{\tilde{\mu}_i^0\}$ and $H_i^0 := X_i \times \{h_i^{*,0}\}$ be the set of belief hierarchies that stop and continue at order 0, respectively. The mth-order belief hierarchies thus specify a signal (in addition to the label $h_i^{*,0}$ or μ_i^0). For reasons that will become clear shortly, it will be convenient to refer to $h_i^{*,0}$ as the (naive) *type*.

We next consider players' beliefs about the state of nature and about whether the other players have stopped reasoning at order 0. Let

$$\widetilde{\Omega}_i^0 := S \times \prod_{j \neq i} \left(\widetilde{H}_j^0 \cup H_j^0 \right);$$

$$\Omega_i^0 := S \times \prod_{j \neq i} H_j^0.$$

---

[6]As is standard, the Cartesian product of a collection of topological spaces $(V_\lambda)_{\lambda \in \Lambda}$ (where Λ is an arbitrary index set) is denoted by V throughout the paper, with typical element v. Given $\lambda \in \Lambda$, we write $V_{-\lambda}$ for $\prod_{\ell \in \Lambda \setminus \{\lambda\}} V_\ell$, with typical element $v_{-\lambda}$.

[7]We thus distinguish between a player's private information or signal and his belief hierarchy or type, as is common in the literature on the robustness of game-theoretic predictions (e.g., Bergemann and Morris, 2005). If a player's signal uniquely determines her belief hierarchy, it is without loss of generality to identify her signal and her belief hierarchy.

[8]The construction can easily be generalized to the case that S and X_i, $i \in N$, are Polish. Assuming that the relevant spaces are compact metric ensures that our results are directly comparable to those of Weinstein and Yildiz (2007).

7

Define the set of first-order belief hierarchies that continue and stop at order 1 by, respectively,

$$\widetilde{H}_i^1 := \widetilde{H}_i^0 \times \Delta\left(\widetilde{\Omega}_i^0\right);$$ (3.1)

$$H_i^1 := \widetilde{H}_i^0 \times \Delta\left(\Omega_i^0\right).$$ (3.2)

Equations (3.1) and (3.2) describe the first-order beliefs for belief hierarchies that reason beyond the first order and that stop reasoning at the first order, respectively, where a first-order belief describes a player's belief about the state of nature and other player's signal. The difference between the belief hierarchies that stop and that continue at order 1 is that the latter can conceive of the possibility that the other players have not yet stopped reasoning at order 0 (i.e., their beliefs are defined on $\prod_{j\neq i} \widetilde{H}_j^0 \cup H_j^0$), while the former can think only that the other players have stopped reasoning at order 0 (i.e., their beliefs are defined on $\prod_{j\neq i} H_j^0$).

For $k = 1, 2, \ldots$, suppose, inductively, that for each player $j \in N$ and all $\ell \leq k$, $\widetilde{H}_j^\ell$ and H_j^ℓ are the sets of belief hierarchies that continue to reason beyond order ℓ and that stop reasoning at that order, respectively. Define

$$\widetilde{H}_i^{\leq k} := \widetilde{H}_i^k \cup \bigcup_{\ell=0}^k H_i^\ell, \qquad \widetilde{\Omega}_i^k := S \times \widetilde{H}_{-i}^{\leq k},$$

$$H_i^{\leq k} := \bigcup_{\ell=0}^k H_i^\ell, \qquad\qquad \Omega_i^k := S \times H_{-i}^{\leq k},$$

and let

$$\widetilde{H}_i^{k+1} := \left\{(x_i, \mu_i^0, \ldots, \mu_i^k, \mu_i^{k+1}) \in \widetilde{H}_i^k \times \Delta(\widetilde{\Omega}_i^k) : \mathrm{marg}_{\widetilde{\Omega}_i^{k-1}} \mu_i^{k+1} = \mu_i^k\right\},$$ (3.3)

$$H_i^{k+1} := \left\{(x_i, \mu_i^0, \ldots, \mu_i^k, \mu_i^{k+1}) \in \widetilde{H}_i^k \times \Delta(\Omega_i^k) : \mathrm{marg}_{\widetilde{\Omega}_i^{k-1}} \mu_i^{k+1} = \mu_i^k\right\}.$$ (3.4)

Again, the interpretation is that $\widetilde{H}_i^{k+1}$ is the set of belief hierarchies that continue to reason at order $k+1$, while the set H_i^{k+1} contains the hierarchies that stop reasoning at $k+1$. As before, the former can conceive of the possibility that the other players have not stopped reasoning at order k, while the latter cannot. A belief hierarchy $h_i^k \in H_i^k$ that stops reasoning at order k is said to have *depth* $d_i(h_i^k) = k$.

The condition on the marginals in Equations (3.3) and (3.4) is a standard coherency condition: it ensures that the beliefs at different orders do not contradict each other (see, e.g., Brandenburger and Dekel, 1993, for a discussion). However, we need to use the extended definition of the marginal (Section 2) here, as we consider the marginal beliefs over belief hierarchies that are still "growing," as well as belief hierarchies that stopped reasoning at some lower order.

In the limit, define

$$H_i^\infty := \left\{(x_i, \mu_i^0, \mu_i^1, \ldots) : (x_i, \mu_i^0, \ldots, \mu_i^k) \in \widetilde{H}_i^k \text{ for all } k \geq 0\right\}.$$

The belief hierarchies in H_i^∞ are those that "reason up to infinity." We therefore say that a belief hierarchy h_i^∞ in H_i^∞ has *infinite depth*, and write $d_i(h_i^\infty) = \infty$.

The next result states that the set H_i^∞ of belief hierarchies with an infinite depth is well-defined:

Lemma 3.1. The set H_i^∞ is nonempty and Polish.

The set H_i^∞ contains the hierarchies with infinite depth, i.e., a hierarchy in H_i^∞ has well-defined beliefs at each order. It will be convenient to define

$$H_i := H_i^\infty \cup \bigcup_{k=0}^{\infty} H_i^k$$

to be the set of all belief hierarchies. Under the usual choice of topology (Section 2), the space H_i is Polish. For future reference, we denote this topology on H_i by τ_i. With some abuse of notation, we sometimes write $(x_i, \mu_i^0, \ldots)$ for a typical element h_i of H_i, regardless of whether the belief hierarchy has a finite or infinite depth; in the former case, there is of course $k < \infty$ such that $h_i = (x_i, \mu_i^0, \ldots, \mu_i^k)$.

3.2   Universal type space

Following Mertens and Zamir (1985), we show that a belief hierarchy not only specifies a belief about other players' higher-order beliefs, but also about their belief hierarchy. We can use that to define a so-called universal type space that contains all type spaces of players with a finite or infinite depth of reasoning, in a sense we make precise.

The first step is to show that a belief hierarchy can be associated with a belief over the set of belief hierarchies of the other players. That is, each belief hierarchy specifies a belief about the full hierarchy of other players, not just about the individual levels of the hierarchy:

Lemma 3.2.

(a) For each belief hierarchy $h_i = (x_i, \mu_i^0, \mu_i^1, \ldots) \in H_i^\infty$ there exists a unique Borel probability measure $\mu_i(h_i)$ on $S \times H_{-i}$ such that

$$\text{marg}_{\widetilde\Omega_i^{\ell-1}} \mu_i(h_i) = \mu_i^\ell$$

for all $\ell = 1, 2, \ldots$.

(b) For each $k > 0$ and every belief hierarchy $h_i = (x_i, \mu_i^0, \mu_i^1, \ldots, \mu_i^k) \in H_i^k$, there exists a unique Borel probability measure $\mu_i(h_i)$ on $S \times H_{-i}^{\leq k-1}$ such that

$$\text{marg}_{\widetilde\Omega_i^{\ell-1}} \mu_i(h_i) = \mu_i^\ell$$

for all $\ell = 1, \ldots, k$.

Thus, each belief hierarchy of player i can be associated with a belief over the set S of states of nature, the signal spaces X_{-i} and over the other players' belief hierarchies, in such a way that i's belief over her ℓth-order space of uncertainty coincides with his ℓth-order belief as specified by her hierarchy of beliefs. That is, the construction is *canonical* in the sense of Brandenburger and Dekel (1993). The result implies that the beliefs of a player at each order she can reason about determine her beliefs about the other players' belief hierarchies. Hence, specifying a player's beliefs about the relevant higher-order spaces of uncertainty fully specifies her beliefs over the full hierarchies of his opponents.

Using Lemma 3.2, we can construct a function that assigns to each belief hierarchy h_i its belief about nature and other players' hierarchies. In the online appendix, we derive the following result:

Corollary 3.3. There is unique mapping $\psi_i : H_i \to \{h_i^{*,0}\} \cup \Delta(S \times H_{-i})$ with the property that for each $k = 1, 2, \ldots$, for each $h_i = (x_i, \mu_i^1, \ldots) \in H_i$ of depth at least k, its kth-order belief μ_i^k is given by

$$\mu_i^k = \mathrm{marg}_{\widetilde{\Omega}_i^{k-1}} \psi_i(h_i),$$

and is such that

- for $h_i^0 \in H_i^0$, $\psi_i(h_i^0) = h_i^{*,0}$;

- for $h_i^k \in H_i^k$, $k < \infty$, the support of $\psi_i(h_i^k)$ lies in $S \times H_{-i}^{\leq k-1}$;

- for $h_i^\infty \in H_i^\infty$, the support of $\psi_i(h_i^\infty)$ lies in $S \times H_{-i}$.

Moreover, the function ψ_i is continuous.

This result is the analogue of the well-known result of Mertens and Zamir (1985) for the space of belief hierarchies of infinite depth.[9] It says that each (nontrivial) belief hierarchy $h_i \in H_i$ for a player $i \in N$ is associated with a belief $\psi_i(h_i)$ on $S \times H_{-i}$.

Mertens and Zamir use the analogue of Corollary 3.3 to construct the so-called *universal Harsanyi type space* $\mathcal{T}^{MZ}$, where the type set of each player is given by a set of belief hierarchies, and the beliefs of each type h_i^{MZ} about the state of nature, signals, and other

---

[9]Unlike the case where all belief hierarchies have an infinite depth of reasoning (Mertens and Zamir, 1985), the function ψ_i is not a homeomorphism. The reason is that there can be belief hierarchies $h_i, h_i' \in H_i$ of different depths such that $\psi_i(h_i) = \psi_i(h_i')$, so that ψ_i is not injective. For example, for $s \in S$ and $x_{-i} \in X_{-i}$ and any $k = 1, 2, \ldots, \infty$, there is a belief hierarchy $h_i^k \in H_i^k$ that assigns probability 1 to $(s, x_{-i}, h_{-i}^{*,0})$. We could rule out certain types that could be deemed "redundant" in this sense (though note that h_i^k and h_i^m have different higher-order beliefs whenever $k \neq m$), and obtain a homeomorphism. This does not affect our main results.

players' types is given by the probability measure $\psi_i^{MZ}(h_i^{MZ})$, where ψ_i^{MZ} is the analogue of the function ψ_i in Corollary 3.3.

In the online appendix, we define the class of type spaces for players with a finite or infinite depth of reasoning, and show that we can likewise construct a *universal type space* $\mathcal{T}^*$ for players with a finite or infinite depth of reasoning. The set of *types* for each player $i \in N$ is given by H_i, and the belief of a type $h_i \in H_i$ about the state of nature, signals, and the other players' types is given by $\psi_i(h_i)$. As in the case of belief hierarchies of infinite depth considered by Mertens and Zamir, the universal type space $\mathcal{T}^*$ generates all belief hierarchies (of finite and infinite depth). With some abuse of terminology, we say that a type has *depth (of reasoning)* $k = 0, 1, \ldots, \infty$ if it generates a belief hierarchy of depth k.

In our analysis, we consider various so-called belief-closed subspaces of the space $\mathcal{T}^*$. A *belief-closed subspace* of $\mathcal{T}^*$ is a tuple $(H_i')_{i \in N}$ such that for $i \in N$, H_i' is a subset of H_i and for each $h_i \in H_i'$ such that $h_i \notin H_i^0$, the support of $\psi_i(h_i)$ lies in $S \times H_{-i}'$.[10] We show in the online appendix that the usual equivalence between type spaces and belief-closed subspaces extend to the current framework: every (nonredundant) type space can be seen as a belief-closed subset of the universal space $\mathcal{T}^*$, and, conversely, any belief-closed subset corresponds to a type space.

The universal Harsanyi type space $\mathcal{T}^{MZ}$ of Mertens and Zamir can be characterized by the event that players have an infinite depth of reasoning, and commonly believe that.

Proposition 3.4. The universal type space $\mathcal{T}^{MZ}$ for Harsanyi type spaces corresponds to the belief-closed subset of $\mathcal{T}^*$ that coincides with the event that all players have an infinite depth of reasoning, and there is common belief that all players have an infinite depth of reasoning.

See the online appendix for a formal statement and a proof. For simplicity, we will say that there is *(correct) common belief in the event that players have an infinite depth of reasoning (CB∞)* when all players have an infinite depth of reasoning, and there is common belief in that event. Henceforth, we will use the terms belief hierarchy and type interchangeably when referring to the elements of H_i, $i \in N$.

4 Games with incomplete information

We define games with incomplete information, and extend the standard concept of rationalizability to our setting. Formally, given a set of players N, set S of states of nature, and signal spaces $(X_i)_{i \in N}$, a *(generalized) Bayesian game* (on $\mathcal{T}^*$) is a tuple $(S, (A_i)_{i \in N}, (u_i)_{i \in N}),$

---

[10]The tuple $(H_i^0)_{i \in N}$ is a belief-closed subset by definition.

where for each player $i \in N$, A_i is i's action set, and $u_i : A \times S \to \mathbb{R}$ is her utility function. Hence, a player's signal does not directly affect her payoff. We assume that payoffs are (jointly) continuous, and that action sets are finite. Players' beliefs are modeled using the universal type space $\mathcal{T}^*$.[11] The set Γ of all games on $A \times S$ is thus the set of profiles $(u_i)_{i \in N}$ of continuous payoff functions on $A \times S$. We endow Γ with its usual sup-norm topology.

It is straightforward to extend the notion of interim correlated rationalizability to this setting (Battigalli and Siniscalchi, 2003; Dekel et al., 2007). In standard Bayesian games this solution concept embodies common belief of rationality; and it allows a type to believe that her opponents' actions are correlated even conditional on them having a particular profile of types and given that a particular state of nature obtains (also see Battigalli et al., 2011). For each player $i \in N$ and $h_i \in H_i \setminus H_i^0$, let

$$R_i^0(h_i) := A_i$$

and, for $k > 1$, define inductively[12]

$$R_i^k(h_i) := \left\{ a_i \in A_i : \begin{array}{l} \text{there is a measurable } \sigma_{-i} : S \times H_{-i} \to \Delta(A_{-i}) \text{ s.t.} \\ \text{supp } \sigma_{-i}(s, h_{-i}) \subseteq \prod_{j \neq i} R_j^{k-1}(h_j) \text{ for all } h_{-i} \in H_{-i}, s \in S; \text{ and} \\ a_i \in \arg\max_{a_i' \in A_i} \int_{S \times H_{-i} \times A_{-i}} u_i(a_i', a_{-i}, s) \sigma_{-i}(s, h_{-i})(a_{-i}) \mathrm{d}\psi_i(h_i) \end{array} \right\}.$$

where supp μ is the support of a probability measure μ. In words, $R_i^k(h_i)$ is the set of actions for type h_i that survive k rounds of iterated deletion of dominated actions: for each action $a_i \in R_i^k(h_i)$, there is a conjecture σ_{-i} that rationalizes it, in the sense that the conjecture has support in the actions of the opponents that have survived $k - 1$ rounds of deletion, and a_i is a best response to this conjecture (given the type's belief $\psi_i(h_i)$).

The *(interim correlated) rationalizable actions* of type h_i for player $i \in N$ are

$$R_i^\infty(h_i) := \bigcap_{k=0}^\infty R_i^k(h_i).$$

If $h_i \in H_i^0$, then we set $R_i^\infty(h_i) := A_i$. That is, any action is rationalizable for a type that doesn't think; this is in the spirit of the assumption that is common in the experimental literature that so-called level-0 types randomize uniformly over their actions (Crawford et al., 2012).

---

[11]In principle, we could have considered a different type space. However, considering the universal type space is without loss of generality in the present context, as the set of interim correlated rationalizable actions for a type depends only on the belief hierarchy that it generates (Dekel et al., 2007).

[12]At first sight, it may seem that our solution concept is not entirely consistent with the idea that players can have a finite depth of reasoning. Specifically, the conjecture σ_{-i} in the definition of $R_i^k(h_i)$ is defined for every type profile $h_{-i} \in H_{-i}$ of i's opponents for every t_i, including type profiles h_{-i}' that correspond to a depth of reasoning that exceeds that of h_i. But since such type profiles h_{-i}' are outside the support of the belief $\psi_i(h_i)$ for type h_i, the beliefs $\sigma_{-i}(s, h_{-i}')$ do not change the surviving set of actions for type h_i.

5   Robust multiplicity

In this section, we show that the seminal "generic uniqueness" result of Weinstein and Yildiz (2007) is not robust to relaxing the assumption that it is common belief that players have an infinite depth of reasoning. We start out by describing the results of Weinstein and Yildiz.

5.1   Generic uniqueness and common belief in an infinite depth

Weinstein and Yildiz prove their generic uniqueness result for the universal type space $\mathcal{T}^{MZ}$ for Harsanyi type spaces constructed by Mertens and Zamir (1985). In the context of the universal space, standard topological notions have a specific meaning, as discussed by Weinstein and Yildiz (p. 372):

- If a subset $O_i \subseteq H_i$ of types for i is *open* in a universal space, then if player i's actual type is in O_i, and our observation of the type's higher-order beliefs is sufficiently precise, then we know that the actual type is in O_i.

- If a subset $V_i \subseteq H_i$ of types for i is *dense* in a universal space, then we cannot rule out the possibility that player i's type lies in V_i, even if we can observe the beliefs of player i up to arbitrarily high order, with arbitrarily high precision.

If a type belongs to an open and dense subset, we say it is *generic*. The complement of an open and dense set has the property that it has an *empty interior*: the complement does not have a nonempty open subset (i.e., the complement of an open and dense subset is nowhere dense). These interpretations are valid both for the universal type space $\mathcal{T}^{MZ}$ for Harsanyi type spaces and the universal type space $\mathcal{T}^*$ for types with a finite or infinite depth.

Weinstein and Yildiz make no a priori restrictions on the domain of payoff structures. Thus, they consider games $(S, (A_i)_{i \in N}, (u_i)_{i \in N})$ whose payoffs are *rich* in the following sense: for every player $i \in N$ and for each action $a_i \in A_i$, there is a state of nature $s^{a_i} \in S$ at which the action a_i is strictly dominant for i. That is,

$$u_i(a_i, a_{-i}, s^{a_i}) > u_i(b_i, a'_{-i}, s^{a_i})$$

for all $a_{-i} \in A_{-i}$ and $b_i \neq a_i$.

We are now ready to state the generic uniqueness result of Weinstein and Yildiz:

Proposition 5.1. (Weinstein and Yildiz, 2007, Prop. 2) The set of types that have a unique rationalizable action in a game with rich payoffs is open and dense in the universal space $\mathcal{T}^{MZ}$ for Harsanyi type spaces. Consequently, the set of types in $\mathcal{T}^{MZ}$ that have multiple rationalizable actions has an empty interior.

This result means that even though there are types in the universal space $\mathcal{T}^{MZ}$ for Harsanyi type spaces that have multiple rationalizable actions, the types in the neighborhood of those types have a unique rationalizable action; we come back to this below. Thus, if we model players' beliefs with a type in $\mathcal{T}^{MZ}$ that has multiple rationalizable actions, then we can conclude that these actions are indeed rationalizable for the player only if we are absolutely certain that the player's actual higher-order beliefs are exactly as given by the type. Hence, multiplicity (of rationalizable actions) is not robust when there is common belief in an infinite depth of reasoning.

5.2   Almost-common belief in an infinite depth

We construct a sequence of sets of types in $\mathcal{T}^*$ that have almost common belief that players have an infinite depth of reasoning. For simplicity, we restrict attention to the case of two players.

We construct a sequence $B_i^1, B_i^2, \ldots$ of types for each player $i = 1, 2$ such that for large n, the types in B_i^n have an infinite depth of reasoning and have almost common belief that players have an infinite depth of reasoning. Formally, for $i = 1, 2$, let

$$B_i^0 := \left\{ h_i \in H_i^\infty : \psi_i(h_i)\Big(\bigcup_{\gamma < \infty} H_{-i}^\gamma \Big) = 1 \right\}$$

be the set of types that have an infinite depth and that believe (with probability 1) that the other player has a finite depth. For $n = 1, 2, \ldots$, define

$$B_i^n := \left\{ h_i \in H_i^\infty : \psi_i(h_i)\Big(B_{-i}^{n-1} \Big) = 1 \right\}.$$

Thus, B_1^1 is the set of types for player 1 with an infinite depth that believe that player 2 has an infinite depth and believes that player 1 has finite depth; the set B_2^1 of types for player 2 is of course defined likewise. Similarly, B_1^2 is the set of types for player 1 with an infinite depth that believe that player 2 has an infinite depth and believes that player 1 has an infinite depth, but believes that player 1 believes that player 2 has a finite depth. Again, B_2^2 is defined similarly. Generally, the types in B_i^n have an infinite depth of reasoning and have nth-order mutual belief in the event that players have an infinite depth of reasoning, but not $(n+1)$th-order mutual belief in that event. When n grows large, the types in B_i^n have *almost common belief in an infinite depth (almost-CB∞)*.

5.3   Generic games

Before stating our results, we define some key notions. We say that there is *robust multiplicity* given the event $E = \prod_{i \in N} E_i \subseteq \prod_{i \in N} H_i$ if for each player $i \in N$, there is a type

$h_i \in E_i$ and an open neighborhood $O_i(h_i)$ of h_i in H_i such that the types in $O_i(h_i)$ have multiple rationalizable actions. If that is the case, multiplicity is not a knife-edge case, unlike the case of CB∞ studied by Weinstein and Yildiz: not only does the type $h_i \in E_i$ have multiple rationalizable actions, the same is true for "nearby" types.

We are particularly interested in the possibility of robust multiplicity under minimal departures of the assumption that it is common belief that there is an infinite depth of reasoning, i.e., in the case where the event E consists of types whose beliefs are arbitrarily close to those of types with common belief in an infinite depth. We say that there is *robust multiplicity under almost common belief in an infinite depth* if for each $n = 0, 1, \ldots$, there is robust multiplicity given $\prod_{i \in N} B_i^n$.

We focus on the class of games studied by Carlsson and van Damme (1993), i.e., two players, two actions games with rich payoffs. Theorem 5.2 says that there is robust multiplicity in generic games in this class, *even if* there is almost-CB∞.[13]

Theorem 5.2. In any generic two player, two action game with rich payoffs, there is robust multiplicity even if there is almost common belief in an infinite depth.

Proof. We establish a sufficient condition on games with rich payoffs under which there is robust multiplicity when there is almost-CB∞; in Appendix C, we show that this condition holds generically.

Fix a game $(S, (A_i)_{i=1,2}, (u_i)_{i=1,2})$ with rich payoffs (and continuous payoff functions) such that each player has two actions. Suppose that for each player i, there exist distinct actions $a_i, b_i \in A_i$ and distinct actions $a_{-i}, b_{-i} \in A_{-i}$ for the other player and $p_i \in [0,1]$ such that

$$p_i u_i(a_i, a_{-i}, s^{a_i}) + (1-p_i)u_i(a_i, a_{-i}, s^{b_i}) > p_i u_i(b_i, a_{-i}, s^{a_i}) + (1-p_i)u_i(b_i, a_{-i}, s^{b_i}), \quad (5.1)$$

and

$$p_i u_i(b_i, b_{-i}, s^{a_i}) + (1-p_i)u_i(b_i, b_{-i}, s^{b_i}) > p_i u_i(a_i, b_{-i}, s^{a_i}) + (1-p_i)u_i(a_i, b_{-i}, s^{b_i}). \quad (5.2)$$

If player $i = 1, 2$ assigns probability p_i to s^{a_i} and probability $1 - p_i$ to s^{b_i}, then, under the conjecture that the other player chooses a_{-i} (regardless of the state of nature), it is a strict best reply for her to play a_i; under the conjecture that the other player chooses b_{-i}, it is a

---

[13]We use the same notion of genericity as in Section 5.1: the set of games for which there is robust multiplicity under almost-CB∞ is open and dense in the set of all games with rich payoffs. If S is finite, the result also holds if we interpret genericity in the measure-theoretic sense: the set of payoff functions with robust multiplicity under almost common belief of infinite depth has probability 1 under the Lebesgue measure on the relevant finite-dimensional space. When S is not finite, there is no obvious measure-theoretic notion that can be applied.

strict best reply for her to choose b_i. If a game satisfies both conditions, then the best reply for each player is *responsive* to the strategy of the other player, at least under certain beliefs about the state of nature. Best replies are responsive in this sense in the game in Figure 3(a): for any $p_i \in (\frac{1}{3}, \frac{2}{3})$ and any type h_i for player i that assigns probability p_i to $s^{a_i} = H$ and $1 - p_i$ to $s^{b_i} = L$, it is a strict best reply for h_i to play a_i under the conjecture that the other player chooses a_i (since $2p_i - (1 - p_i) > 0$), while it is a strict best reply for h_i to choose b_i under the conjecture that the other player chooses b_i (since $p_i - 2(1 - p_i) < 0$). By contrast, in the game in Figure 3(b), best replies do not depend on conjectures about the other's strategy, only on beliefs about the state of nature. We refer to games with responsive best replies as *responsive games*, and to games that do not have responsive best replies as *nonresponsive games*.

	a_2	b_2
a_1	2,2	1,0
b_1	0,1	0,0

$s = H$

	a_2	b_2
a_1	-1,-1	-2,0
b_1	0,-2	0,0

$s = L$

(a)

	a_2	b_2
a_1	1,1	1,0
b_1	0,1	0,0

$s = H$

	a_2	b_2
a_1	0,0	0,1
b_1	1,0	1,1

$s = L$

(b)

Figure 3: (a) A game where players' best replies are responsive; (b) A game where players' best replies are not responsive.

We show that in responsive games, for each player $i = 1, 2$, $n = 0, 1, \ldots$, there is $h_i \in B_i^n$ and an open neighborhood of h_i in H_i such that both actions are rationalizable for the types in the neighborhood.

To prove this, we first show that for each player $i = 1, 2$, there is a type of finite depth and an open neighborhood of that type such that both actions are rationalizable for the types in the neighborhood. It will be useful to introduce some unifying notation. For $i = 1, 2$ and $\alpha = 0, 1, \ldots$, let $M_i^\alpha := H_i^\alpha$ be the set of types of depth α. To define the sets B_i^n, $n = 0, 1, \ldots$, we need to consider the transfinite ordinals $\omega, \omega + 1, \ldots$.[14] Let $M_i^\omega := B_i^0$, and for $n = 1, 2, \ldots$, let $M_i^{\omega + n} := B_i^n$. We show that for any set M_i^α, $\alpha < \omega + \omega$, there is a type $h_i^\alpha \in M_i^\alpha$ and a neighborhood $O_i^\alpha(h_i^\alpha)$ of h_i^α in H_i such that both actions are rationalizable for the types in $O_i^\alpha(h_i^\alpha)$.

The claim follows immediately for $\alpha = 0$, as both actions are rationalizable for any type $h_i^0 \in H_i^0$, $i = 1, 2$, and $O_i^0(h_i^0) := H_i^0$ is an open neighborhood for any such type. For $\alpha > 0$, suppose that for each player $i = 1, 2$ and $\gamma < \alpha$, there is a type $h_i^\gamma \in M_i^\gamma$ and an open neighborhood $O_i^\gamma(h_i^\gamma)$ of h_i^γ such that both actions are rationalizable for the types in $O_i^\gamma(h_i^\gamma)$.

---

[14] Recall that $\omega = \{0, 1, \ldots\}$ is the first infinite ordinal, and that $\omega + n$, $n = 1, 2, \ldots$, is the successor of $\omega + n - 1$. The ordinal $\omega + \omega$ is the first ordinal that is greater than $0, 1, \ldots, \omega, \omega + 1, \omega + 2, \ldots$.

Suppose that α is finite or that $\alpha = \omega$. Let $h_i^\alpha \in M_i^\alpha$ be a type such that

$$\psi_i(h_i^\alpha)\left(s^{a_i}, \bigcup_{\gamma < \alpha} O_{-i}^\gamma(h_{-i}^\gamma)\right) = p_i, \qquad \psi_i(h_i^\alpha)\left(s^{b_i}, \bigcup_{\gamma < \alpha} O_{-i}^\gamma(h_{-i}^\gamma)\right) = 1 - p_i.$$

Since (5.1) and (5.2) hold with strict inequality, there is $\eta > 0$ such that for each player i,

$$p_i\big(u_i(a_i, a_{-i}, s^{a_i}) - \eta\big) + (1 - p_i)\big(u_i(a_i, a_{-i}, s^{b_i}) - \eta\big) > p_i\big(u_i(b_i, a_{-i}, s^{a_i}) + \eta\big) + (1 - p_i)\big(u_i(b_i, a_{-i}, s^{b_i}) + \eta\big),$$

and

$$p_i\big(u_i(b_i, b_{-i}, s^{a_i}) - \eta\big) + (1 - p_i)\big(u_i(b_i, b_{-i}, s^{b_i}) - \eta\big) > p_i\big(u_i(a_i, b_{-i}, s^{a_i}) + \eta\big) + (1 - p_i)\big(u_i(a_i, b_{-i}, s^{b_i}) + \eta\big),$$

where $a_i, b_i, a_{-i}, b_{-i}, s^{a_i}, s^{b_i}$ are as in (5.1) and (5.2). Let d_S be a metric that metrizes the topology on S. By continuity of the payoff functions, there exists $\zeta > 0$ (depending on η) such that for all $\tilde{s}^{a_i}, \tilde{s}^{b_i}$ such that $d_S(\tilde{s}^{a_i}, s^{a_i}) < \zeta$ and $d_S(\tilde{s}^{b_i}, s^{b_i}) < \zeta$,

$$
\begin{aligned}
&u_i(a_i, a_{-i}, \tilde{s}^{a_i}) > u_i(a_i, a_{-i}, s^{a_i}) - \eta, && u_i(a_i, a_{-i}, \tilde{s}^{b_i}) > u_i(a_i, a_{-i}, s^{b_i}) - \eta, \\
&u_i(b_i, a_{-i}, \tilde{s}^{a_i}) < u_i(b_i, a_{-i}, s^{a_i}) + \eta, && u_i(b_i, a_{-i}, \tilde{s}^{b_i}) < u_i(b_i, a_{-i}, s^{b_i}) + \eta, \\
&u_i(b_i, b_{-i}, \tilde{s}^{a_i}) > u_i(b_i, b_{-i}, s^{a_i}) - \eta, && u_i(b_i, b_{-i}, \tilde{s}^{b_i}) > u_i(b_i, b_{-i}, s^{b_i}) - \eta, \\
&u_i(a_i, b_{-i}, \tilde{s}^{a_i}) < u_i(a_i, b_{-i}, s^{a_i}) + \eta, && u_i(a_i, b_{-i}, \tilde{s}^{b_i}) < u_i(a_i, b_{-i}, s^{b_i}) + \eta.
\end{aligned}
$$

Let $V_{a_i}(\zeta)$ and $V_{b_i}(\zeta)$ be the open ζ-balls around s^{a_i} and s^{b_i}, respectively, that is, let $V_{a_i}(\zeta) := \{\tilde{s}^{a_i} \in S : d_S(\tilde{s}^{a_i}, s^{a_i}) < \zeta\}$, and let $V_{b_i}(\zeta)$ be defined similarly. Also, write $O_{-i}^{<\alpha}$ for $\bigcup_{\gamma < \alpha} O_{-i}^\gamma(h_{-i}^\gamma)$. Then, for $\xi > 0$, define

$$
\begin{aligned}
U_i^\xi(h_i^\alpha) := &\big\{h_i \in H_i : \psi_i(h_i)\big(V_{a_i}(\zeta) \times O_{-i}^{<\alpha}\big) > \psi_i(h_i^\alpha)\big(V_{a_i}(\zeta) \times O_{-i}^{<\alpha}\big) - \xi\big\} \cap \\
&\big\{h_i \in H_i : \psi_i(h_i)\big(V_{b_i}(\zeta) \times O_{-i}^{<\alpha}\big) > \psi_i(h_i^\alpha)\big(V_{b_i}(\zeta) \times O_{-i}^{<\alpha}\big) - \xi\big\}.
\end{aligned}
$$

By continuity of the belief function ψ_i, the set $U_i^\xi(h_i^\alpha)$ is open in H_i (Billingsley, 1968, App. III). Moreover, it contains h_i^α.

Since the payoff function u_i is defined on a compact space, there is $M \in [0, \infty)$ such that $u_i(\tilde{a}, \tilde{s}) \in [-M, M]$ for all $(\tilde{a}, \tilde{s}) \in A \times S$. It follows that for every type $\xi > 0$ and $h_i \in U_i^\xi(h_i^\alpha)$, under the conjecture that the other player plays a_{-i}, the expected payoff to playing a_i is at least

$$(p_i - \xi)\big(u_i(a_i, a_{-i}, s^{a_i}) - \eta\big) + (1 - p_i - \xi)\big(u_i(a_i, a_{-i}, s^{b_i}) - \eta\big) - 2\xi(M + \eta),$$

while the expected payoff of playing b_i is at most

$$(p_i - \xi)\big(u_i(b_i, a_{-i}, s^{a_i}) + \eta\big) + (1 - p_i - \xi)\big(u_i(b_i, a_{-i}, s^{b_i}) + \eta\big) + 2\xi(M + \eta).$$

17

Consequently, there is $\bar{\xi}_1 > 0$ such that for all $\xi < \bar{\xi}_1$, the former expression is at least as great as the latter, and a_i is rationalizable for any $h_i \in U_i^\xi(h_i^\alpha)$. Likewise, for any $\xi > 0$ and type $h_i \in U_i^\xi(h_i^\alpha)$, under the conjecture that the other player plays b_{-i}, the expected payoff to playing b_i is at least

$$(p_i - \xi)\big(u_i(b_i, b_{-i}, s^{a_i}) - \eta\big) + (1 - p_i - \xi)\big(u_i(b_i, b_{-i}, s^{b_i}) - \eta\big) - 2\xi M,$$

while the expected payoff of playing b_i is at most

$$(p_i - \xi)\big(u_i(a_i, b_{-i}, s^{a_i}) + \eta\big) + (1 - p_i - \xi)\big(u_i(a_i, b_{-i}, s^{b_i}) + \eta\big) + 2\xi M.$$

Again, there is $\bar{\xi}_2 > 0$ such that for any $\xi < \bar{\xi}_2$, action b_i is rationalizable for any $h_i \in U_i^\xi(h_i^\alpha)$. If we define $\bar{\xi} := \min\{\bar{\xi}_1, \bar{\xi}_2\}$, the result now follows by setting $O_i^\alpha(h_i^\alpha) := U_i^\xi(h_i^\alpha)$ for some $\xi \in (0, \bar{\xi})$. The proof for the case that $\alpha = \omega + n$ for some $n = 1, 2, \dots$ is similar and therefore omitted.

	a_2	b_2			a_2	b_2
a_1	$1 + \eta_1, 1 + \eta_2$	$1 + \eta_3, \eta_4$		a_1	η_9, η_{10}	$\eta_{11}, 1 + \eta_{12}$
b_1	$\eta_5, 1 + \eta_6$	η_7, η_8		b_1	$1 + \eta_{13}, \eta_{14}$	$1 + \eta_{15}, 1 + \eta_{16}$
	$s = H$				$s = L$	

(c)

Figure 4: A perturbation of the game in Figure 3(b). Under generic payoff perturbations $\eta_1, \dots, \eta_{16}$, the perturbed game is responsive.

In the appendix, we show that generically, games are responsive. Intuitively, if we perturb the payoffs of a game with responsive best replies by a little bit, best replies still depend on conjectures about the other player's action. Hence, the set of responsive games is open in the set of games with rich payoffs. On the other hand, if we perturb the payoffs of a nonresponsive game slightly, then a player's best reply may depend on her conjecture about the other's play for at least some beliefs about the state of nature, unless we choose the perturbations in a very particular way. For example, consider the game in Figure 4. This is a perturbation of the game in Figure 3(b). For "most" (small) values of $\eta_1, \dots, \eta_{16}$, Eqs. (5.1) and (5.2) are satisfied for p_i close to $\frac{1}{2}$. This means that the set of responsive games is dense in the set of games with rich payoffs. $\square$

The condition that there is almost common belief in an infinite depth is sufficient for there to be robust multiplicity, but obviously not necessary. As the proof makes clear, there is also robust multiplicity when players have a finite depth of reasoning. The strength of the

result is precisely that we get robust multiplicity *even if* we relax the assumption that there is common belief in an infinite depth only a little bit. As we discuss in Section 6, there are other interesting types that exhibit robust multiplicity besides the ones that satisfy almost-CB∞.

Theorem 5.2 implies that the result of Weinstein and Yildiz that types with CB∞ generically have a unique rationalizable action in games with rich payoffs does not extend if we relax the assumption of common belief in an infinite depth of reasoning:[15]

Corollary 5.3. Fix any generic (two player, two action game) with rich payoffs. For every player $i = 1, 2$, $n = 0, 1, \ldots$, the set of types in B_i^n with multiple rationalizable actions has a nonempty open subset in $\mathcal{T}^*$.

Theorem 5.2 and Corollary 5.3 allow for general games, and do not put any a priori restrictions on players' beliefs about payoffs. However, in many applications, the information structure imposes natural restrictions on players' beliefs. In the next section, we show that there is robust multiplicity even for types whose beliefs are consistent with the information structure.

5.4 Global games

We prove our result in the context of one of the seminal games of the literature on global games (Carlsson and van Damme, 1993; Morris and Shin, 2003), which we also discussed in Section 1. Two players, labeled by $i = 1, 2$, decide simultaneously whether to invest (I) or not to invest (NI). The payoff matrix is

	I	NI
I	s, s	$s-1, 0$
NI	$0, s-1$	$0, 0$

The set of states of nature is $S := [-1, 2]$. If $s > 1$ investing (I) is strictly dominant for each player. If $s < 0$, then not investing (NI) is strictly dominant for each player. For intermediate values of the state s, i.e., $s \in [0, 1]$, both actions are rationalizable.

The information structure is as follows. The state s is drawn from S according to the uniform distribution. Prior to playing, each player i receives a signal $x_i \in X_i := [-1, 2]$ that is informative about the state of nature. When state of nature s, players' signals are drawn independently from the interval $[s - \varepsilon, s + \varepsilon]$, where $\varepsilon > 0$ is small ($\varepsilon < \frac{1}{2}$, say). Given the

---

[15] The result of Weinstein and Yildiz (2007) applies to games (with rich payoffs) with an arbitrary finite set of players and finite action sets. To show that their result does not carry over, it is sufficient to show that it does not hold for some games. That said, we could show a weaker result for games with an arbitrary (finite) number of players and actions if we relax the assumption that there is almost common belief in infinite depth.

signal x_i, player i's posterior on the state of nature is uniform in $[x_i - \varepsilon, x_i + \varepsilon] \cap [-1, 2]$ (by Bayes' rule). Moreover, conditional on $s \in [x_i - \varepsilon, x_i + \varepsilon] \cap [-1, 2]$, the player believes that the other player's signal x_{-i} is distributed uniformly on $[s - \varepsilon, s + \varepsilon]$.

As shown by Carlsson and van Damme, if all types have an infinite depth of reasoning and this is commonly believed, then I is the only rationalizable action for a player who have received a signal $x_i < \frac{1}{2}$, and NI is the unique rationalizable action for a player who received a signal $x_i > \frac{1}{2}$. Types that have received signal $x_i = \frac{1}{2}$ have multiple rationalizable actions: if such a type believes that the other player chooses his dominant action (i.e., invests whenever his signal is greater than $\frac{1}{2}$ and does not invest otherwise), then the type is indifferent between I and NI. By the generic uniqueness result of Weinstein and Yildiz (Proposition 5.1), this multiplicity is not robust: if the type's beliefs are perturbed slightly, then both actions are rationalizable.

We next show that that there is robust multiplicity under this information structure. To show this, we define the maximal belief-closed subset of $\mathcal{T}^*$ that is consistent with the information structure and in which types can have a finite or infinite depth of reasoning. More precisely, for $\varepsilon > 0$, we define the belief-closed subset $(T_i^\varepsilon)_{i=1,2}$ of $\mathcal{T}^*$, with type set T_i^ε for player $i = 1, 2$, that has the following properties:

(1) the belief of each type about the state of nature and the other player's signal is as described above;

(2) the belief of each type about the other's signal is independent of its belief about the other's depth of reasoning;

(3) $(T_i^\varepsilon)_{i=1,2}$ is the maximal belief-closed subspace that is consistent with the information structure in the sense that it contains any belief-closed subspace that satisfies the above two conditions;

see Appendix A for a formal definition. We denote this belief-closed subset by $\mathcal{T}^\varepsilon$. The assumption that players' beliefs about signals are independent from their beliefs about players' depth of reasoning seems a natural one; and we focus on the maximal belief-closed subspace to obtain the strongest possible result.

Appendix B defines a collection $\mathcal{C}_i$, $i = 1, 2$, of (equivalence) classes of belief hierarchies for player i that partitions H_i and is such that each class $C_i \in \mathcal{C}_i$ is characterized completely by its beliefs about players' depth of reasoning. If the class $C_i \in \mathcal{C}_i$ is included in the set B_i^n of types with nth-order mutual belief in the event that players have an infinite depth, then we say that C_i *satisfies nth-order mutual belief in infinite depth*. These equivalence classes are critical for obtaining the strongest possible result, as we discuss below.

Theorem 5.4 shows that even if there is almost common belief that players have an infinite depth of reasoning, and even if a player has received a signal arbitrarily close to the dominance regions (i.e., x_i close to 0 or close to 1), if there is some grain of naiveté and signals are sufficiently precise, then both actions are rationalizable, and this multiplicity is robust to small perturbations in beliefs.

Theorem 5.4. For each player $i = 1, 2$, signal $x_i \in (0, 1)$, $n = 0, 1, \ldots$, and class $C_i^n \in \mathcal{C}_i$ of beliefs that satisfies nth-order mutual belief in infinite depth, there is $\bar{\varepsilon} > 0$ such that for all $\varepsilon < \bar{\varepsilon}$, every type $h_i \in T_i^\varepsilon$ with signal x_i and with beliefs in class C_i^n has an open neighborhood in H_i such that both actions are rationalizable for the types in this neighborhood.

The proof is relegated to Appendix C. The challenge in proving this result is that, as we move closer to the dominance regions and consider higher and higher orders of mutual belief in infinite depth, multiplicity is possible only if the noise level ε is sufficiently small.

The proof demonstrates that we can nevertheless bound the upper bound $\bar{\varepsilon}$ on the noise level away from zero even for signals close to the dominance regions and when there is almost common belief in infinite depth. This is possible, because, as we show, no matter what a type h_i believes about the signal x_{-i} of its opponent, and what order of mutual belief in infinite depth it satisfies, we can always find finitely many types of its opponents to which it assigns a sufficiently high probability and that have multiple rationalizable actions for some upper bound on the noise that is greater than 0. This allows us to bound the upper bound for type h_i in turn.

To show this, we exploit the assumption that players' beliefs about signals are independent from their beliefs about players' depth of reasoning. The independence assumption allows us to choose the lower bound so that it holds uniformly for a class C_i of types that are characterized solely by their beliefs about players' depth. This makes that our result holds for all types in T_i^ε with beliefs in that class, even if they put a very large weight on types that are only slightly less sophisticated than they are.

By working with the equivalence classes C_i, we can choose the bound on the noise level *after* fixing the signal and the class (and thus the order of mutual belief in infinite depth), but *before* choosing the type. This gives us the strongest possible result: the bound does not depend on the specifics of the belief hierarchies, only on the two main dimensions (the signal and beliefs about depth of reasoning); it is easy to see that there is no bound that would hold uniformly across types. While we prove Theorem 5.4 for this particular game and information structure, we believe the arguments can be applied more generally.

Together, Theorems 5.2 and 5.4 show that there are many types, possibly with almost common belief in an infinite depth of reasoning, that have multiple rationalizable actions in

this game, and this multiplicity is robust to further perturbations of beliefs. On the other hand, the results of Weinstein and Yildiz (2007) imply that types that have common belief in an infinite depth of reasoning generically have a unique rationalizable action in games with rich payoffs. Hence, both uniqueness and multiplicity are common when we allow types to have any depth of reasoning.

6 Discussion and related literature

Robust multiplicity given other events   We have focused on showing that we have robust multiplicity for the types in B_i^n, $i = 1, 2$, $n = 0, 1, \ldots$, so as to show that the generic uniqueness result of Weinstein and Yildiz does not hold if we slightly relax the assumption that there is common belief in an infinite depth of reasoning. However, there are other interesting sets of types for which one can show similar results. For example, we could consider a sequence of types that are very sophisticated (in the sense that they have a high (but finite) or even infinite depth of reasoning) and assign an arbitrarily high probability to the event that the other player is only slightly less sophisticated.

Formally, we construct a sequence of sets $\{H_i^{\alpha,\eta}\}_\alpha$, where $\eta > 0$ and α is a countable ordinal.[16] For $i = 1, 2$, let $H_i^{0,\eta} := X_i \times \{h_i^{*,0}\}$. For $\alpha > 0$, assume, inductively, that the sets $H_i^{\gamma,\eta}$ have been defined for $i = 1, 2$ and $\gamma < \alpha$. We distinguish between the case where α is a successor ordinal and where it is a limit ordinal.[17] If α is a successor ordinal (e.g., $\alpha = 1, 2, \ldots$), then define

$$H_i^{\alpha,\eta} := \left\{ h_i \in H_i \setminus \bigcup_{\gamma < \alpha} H_i^{\gamma,\eta} : \psi_i(h_i)\left(\bigcup_{\gamma < \alpha} H_{-i}^{\gamma,\eta}\right) = 1, \psi_i(h_i)\left(H_{-i}^{\alpha-1,\eta}\right) > 1 - \eta \right\}.$$

to be the set of types that assigns probability greater than $1 - \eta$ to the types in $H_{-i}^{\alpha-1}$. These types assign a high probability (more than $1 - \eta$) to the event that the other is only slightly less sophisticated than they are (i.e., to $H_{-i}^{\alpha-1,\eta}$). If α is a limit ordinal, then there is no unique largest ordinal smaller than α (or even a finite set of such ordinals). So, to model that a type at level α assigns high probability to types that are of slightly lower level, we fix an increasing sequence $\{\alpha_n\}_{n=1}^\infty$ of ordinals converging to α, and require that types assign high probability

---

[16]The result does not extend immediately to the uncountable ordinals, as not all relevant sets may be measurable in that case.

[17]Recall that an ordinal α can be identified with the set $\{\gamma : \gamma < \alpha\}$ of its predecessors; we identify the finite ordinals with the natural numbers $0, 1, 2, \ldots$, so that the first infinite ordinal ω is equal to $\{0, 1, \ldots\}$. An ordinal is a *successor ordinal* if it is the successor of some ordinal, where the successor of an ordinal α is the least ordinal greater than α. An ordinal is a *limit ordinal* if it is not 0 or a successor ordinal. For example, 37 is the successor of 36, and ω is a limit ordinal.

to the tail of that sequence.[18] Then, define

$$H_i^{\alpha,\eta} := \left\{ h_i \in H_i \setminus \bigcup_{\gamma < \alpha} H_i^{\gamma,\eta} : \psi_i(h_i)\left(\bigcup_{\gamma < \alpha} H_{-i}^{\gamma,\eta}\right) = 1, \psi_i(h_i)\left(\bigcup_{n > \frac{1}{\eta}} H_{-i}^{\alpha_n,\eta}\right) > 1 - \eta \right\}.$$

Thus, the types in this set assign a high probability (more than $1 - \eta$) to types of the other player that are only slightly less sophisticated ($n > \frac{1}{\eta}$).

We then have the following analogues to Theorems 5.2 and 5.4:

Theorem 6.1. In any generic two player, two action game with rich payoffs, for any $\eta > 0$, player $i = 1, 2$, countable ordinal α, there is a type $h_i^{\alpha,\eta} \in H_i^{\alpha,\eta}$ and an open neighborhood of $h_i^{\alpha,\eta}$ such that both actions are rationalizable for the types in the neighborhood.

Theorem 6.2. For each $\eta > 0$, player $i = 1, 2$, signal $x_i \in (0, 1)$, class $C_i^{\alpha,\eta} \in \mathcal{C}_i$ of beliefs such that $C_i^{\alpha,\eta} \subseteq H_i^{\alpha,\eta}$, there is $\bar{\varepsilon} > 0$ such that for all $\varepsilon < \bar{\varepsilon}$, every type $h_i \in T_i^\varepsilon$ with signal x_i and with beliefs in class $C_i^{\alpha,\eta}$ has an open neighborhood in H_i such that both actions are rationalizable for the types in this neighborhood.

The proof are essentially identical to those of Theorems 5.2 and 5.4, respectively, and are therefore omitted.[19]

Minimum depth k   Our results do not rely on the fact that the space of belief hierarchies includes belief hierarchies (types) in H_i^0 that do not have beliefs and for whom any action is rationalizable. We could have constructed the set of belief hierarchies in which every belief hierarchy has at least depth $k < \infty$, and our results would go through with minor modifications, as long as there are types with common belief in an infinite depth of reasoning for which multiple actions survive k rounds of elimination of dominated strategies. These types can be used to construct an open set of types with depth k for which multiple actions are rationalizable. The rest of the proofs of Theorems 5.2 and 5.4 then goes through without change.

---

[18]If α is less than the first the least critical ordinal $\epsilon_0 = \omega^{\omega^{\omega^{\cdot^{\cdot^{\cdot}}}}}$, there is a standard such sequence (e.g., Takeuti, 1987, p. 120); for larger ordinals, we can fix an arbitrary such sequence.

[19]Theorems 5.2 and 6.1 do not imply each other, and likewise for Theorems 5.4 and 6.2: for α finite, $H_i^{\alpha,\eta}$ is neither a subset or a superset of the set H_i^α of types of depth α. Even if we make $H_i^{\alpha,\eta}$ a subset of H_i^α for α finite by requiring that the types in $H_i^{\alpha,\eta}$ have depth α (and that the types in $H_i^{\alpha,\eta}$ for infinite α have infinite depth), then still $H_i^{\omega+1,\eta}$ is not a subset of B_i^1 or vice versa. Also, while Theorem 5.2 and Theorem 5.4 consider only the sets B_i^α for finite ordinals α, we could have considered all countable ordinals, though the intuitive interpretation of the set B_i^α for infinite α is not immediate.

Topology   While we have used the standard topology in our construction of the set of belief hierarchies (Section 2), a natural question is whether there are alternative topologies that one could use, and, if so, whether results are sensitive to this choice. An alternative construction, which leads to a slightly weaker topology than the current one, is considered by Kets (2010, 2013). The main difference is that in this alternative topology, the set of types of infinite depth is no longer open.[20] However, this does not affect our main results: what we need for our results to hold is that there is a set of types of finite depth $k \geq 0$ that is open. To the best of our knowledge, this holds for all topologies that coincide with the standard product topology on the set of belief hierarchies for a given depth of reasoning. Indeed, under any natural topology, types of different depths are, in fact, different.

Common p-belief in infinite depth   In this paper, we use the notion of almost common belief as employed by Rubinstein (1989) and Carlsson and van Damme (1993): a belief hierarchy has almost common belief in an event E if it has nth-order mutual belief in E for high n. Monderer and Samet (1989) have defined a different notion of almost common belief than the one we use here: an event E is almost common belief in their sense if there is p close to 1 such that all players assign probability at least p to E, assign probability at least p to the event that all players assign probability at least p to E, and so on.

Our results do not go through under this alternative notion of almost common belief. This is not surprising, since strategic behavior is generally sensitive to the notion of almost common belief used (e.g., Monderer and Samet, 1989). More specifically, the proofs of Theorems 5.2 and 5.4 use the fact that we can always construct an open set of types that have multiple rationalizable actions to which a type assigns high probability even if it has almost-CB∞ in our sense. By the generic uniqueness result of Weinstein and Yildiz (2007), such an open set of types does not exist under the alternative notion of almost-CB∞. Just like the results of Weinstein and Yildiz, Carlsson and van Damme (1993), and others require that there is a "seed" of uniqueness, by there being an open set of types with a unique rationalizable action, our results rely on the existence of a similar seed of multiplicity.

Related literature   A number of papers have shown that the results of Weinstein and Yildiz do not hold if the topology is changed (Monderer and Samet, 1989; Dekel et al., 2006; Chen et al., 2010) or when results are required to be robust against small misspecifications of the game (Chen et al., 2013). Unlike these papers, we do not change the topology or introduce additional uncertainty about the game. Rather, we identify a plausible epistemic assumption

---

[20]This implies in particular, that a sequence $h_i^1, h_i^2, \ldots$ of belief hierarchies of (strictly) increasing depth can converge to a belief hierarchy of infinite depth.

– that it need not be commonly known that players have an infinite depth of reasoning –, and investigate the consequences.[21]

Strzalecki (2009) likewise studies the behavior of players with a finite depth of reasoning. He shows that types with a finite depth of reasoning can attack in equilibrium in the e-mail game of Rubinstein (1989) if they exchange enough messages, which is ruled out in the case that players have an infinite depth of reasoning. However, Strzalecki does not address the issue of multiplicity and whether multiplicity can be robust. These questions cannot be addressed in his framework, as it does not model players' beliefs about the state of nature. Moreover, the set of belief hierarchies constructed by Strzalecki does not include the belief hierarchies with an infinite depth of reasoning. This means that it is not possible in his framework to study outcomes under almost common belief in an infinite depth of reasoning.[22]

Finally, Kets (2010, 2013) presents an even richer model of the higher-order beliefs of players with a potentially finite depth than the one used here. Similar results to the ones presented here would hold in the framework of Kets, so that our results are robust to the specification of higher-order beliefs. Since we do not need the additional richness of the framework of Kets, we use the present, simpler, framework.

Appendix A   The space $\mathcal{T}^{\varepsilon}$

Let $\varepsilon > 0$. We define a belief-closed subspace $\mathcal{T}^{\varepsilon}$ of the universal space $\mathcal{T}^{*}$ that satisfies the following properties:

(1) The belief of each type about S and X_{-i} is consistent with the information structure. That is, if type $t_i \in T_i^{\varepsilon}$ has depth greater than 0 and has observed signal x_i, it assigns probability 1 to the event that the state of nature is uniform in $[x_i - \varepsilon, x_i + \varepsilon] \cap [-1, 2]$, and that conditional on $s \in [x_i - \varepsilon, x_i + \varepsilon] \cap [-1, 2]$, the other player's signal x_{-i} is distributed uniformly on $[s - \varepsilon, s + \varepsilon]$.

---

[21]Other potential sources of multiplicity in coordination games with incomplete information include endogenous information about the state of nature, signaling, and learning; see, e.g., Angeletos and Werning (2006), Angeletos et al. (2006, 2007), and Yang (2013). Unlike these models, we obtain multiplicity even if the game is static, and there is no learning, signaling, or endogenous information.

[22]In the online appendix, we show that the universal type space $\mathcal{T}^{*}$ (strictly) contains the set of belief hierarchies constructed by Strzalecki (2009). As we discuss there, the universal type space $\mathcal{T}^{*}$ constructed here contains many more belief hierarchies than just those of finite depth and the belief hierarchies in the universal type space of Mertens and Zamir (1985), so that the current framework is much richer than that of Strzalecki.

(2) The belief of each type about the other's signal is independent of its belief about the other's depth of reasoning.

(3) Any belief-closed subspace that satisfies the above two properties is contained in $\mathcal{T}^\varepsilon$.

A.1   Preliminaries

To formally define this space, some more notation will be useful. We first define auxiliary topologies on the spaces of belief hierarchies defined in Section 3 that distinguish belief hierarchies only on the basis of their beliefs about players' depth of reasoning. This will allow us to formally express the independence condition (2) above.

Formally, for $i = 1, 2$, let $\tau_{\widetilde{H}_i^0}$ be the trivial topology on $\widetilde{H}_i^0$, i.e., $\tau_{\widetilde{H}_i^0} = \{\widetilde{H}_i^0, \emptyset\}$; likewise, let $\tau_{H_i^0}$ be the trivial topology on H_i^0. Let $\tau_{\Delta(\widetilde{\Omega}_i^0)}$ be the topology on $\Delta(\widetilde{\Omega}_i^0)$ generated by the sets

$$\{\mu_i \in \Delta(\widetilde{\Omega}_i^0) : \mu_i(S \times G_{-i}) > \nu_i(S \times G_{-i}) - \delta\} : \qquad \nu_i \in \Delta(\widetilde{\Omega}_i^0), G_{-i} \in \tau_{\widetilde{H}_{-i}^0}, \delta > 0;$$

and let the topology $\tau_{\Delta(\Omega_i^0)}$ on $\Delta(\Omega_i^0)$ be defined analogously.[23] Let $\tau_{\widetilde{H}_i^1}$ and $\tau_{H_i^1}$ be the product topologies on $\widetilde{H}_i^1$ and H_i^1, respectively, induced by these topologies.

For $k \geq 1$, suppose that for all $\ell \leq k$, the topologies $\tau_{\widetilde{H}_i^\ell}$ and $\tau_{H_i^\ell}$ on $\widetilde{H}_i^\ell$ and H_i^ℓ have been defined, and let $\tau_{\widetilde{H}_i^{\leq k}}$ and $\tau_{H_i^{\leq k}}$ be the sum topologies on $\widetilde{H}_i^{\leq k}$ and $H_i^{\leq k}$, respectively. Let $\tau_{\Delta(\widetilde{\Omega}_i^k)}$ be the topology on $\Delta(\widetilde{\Omega}_i^k)$ generated by the sets

$$\{\mu_i \in \Delta(\widetilde{\Omega}_i^k) : \mu_i(S \times G_{-i}) > \nu_i(S \times G_{-i}) - \delta\} : \qquad \nu_i \in \Delta(\widetilde{\Omega}_i^k), G_{-i} \in \tau_{\widetilde{H}_{-i}^{\leq k}}, \delta > 0;$$

and let the topology $\tau_{\Delta(\Omega_i^k)}$ on $\Delta(\Omega_i^k)$ be defined similarly.

Let $\tau_{\widetilde{H}_i^{k+1}}$ be the relative topology on $\widetilde{H}_i^{k+1}$ induced by the product topology induced by $(\widetilde{H}_i^k, \tau_{\widetilde{H}_i^k})$ and $(\Delta(\widetilde{\Omega}_i^k), \tau_{\Delta(\widetilde{\Omega}_i^k)})$; let the topology $\tau_{H_i^{k+1}}$ on H_i^{k+1} be defined similarly. Let $\tau_{H_i^\infty}$ be the relative topology on H_i^∞ induced by the product topology generated by $(\widetilde{H}_i^0, \tau_{\widetilde{H}_i^0})$, $(\Delta(\widetilde{\Omega}_i^0), \tau_{\Delta(\widetilde{\Omega}_i^0)})$, $(\Delta(\widetilde{\Omega}_i^1), \tau_{\Delta(\widetilde{\Omega}_i^1)})$, ..., and let τ_{H_i} be the sum topology on H_i induced by $\tau_{H_i^0}$, $\tau_{H_i^1}$, ..., $\tau_{H_i^\infty}$; denote the Borel σ-algebra generated by τ_{H_i} by $\mathscr{B}_{H_i}$. Note that every nonempty open set in τ_{H_i} (or: Borel set) is of the product form. That is, if $U_i \in \tau_{H_i}$ is nonempty, then there exist $G_i \subseteq H_i$ such that $U_i = X_i \times \{(\mu_i^0, \ldots) : (x_i, \mu_i^0, \ldots) \in G_i \text{ for some } x_i \in X_i\}$. Finally, let $\tau_{\Delta(S \times H_{-i})}$ be the topology on $\Delta(S \times H_{-i})$ generated by the sets

$$\{\mu_i \in \Delta(S \times H_{-i}) : \mu_i(S \times G_{-i}) > \nu_i(S \times G_{-i}) - \delta\} \qquad G_{-i} \in \tau_{H_{-i}}.$$

---

[23]We emphasize that the spaces $\widetilde{H}_i^0$, H_i^0, $\Delta(\widetilde{\Omega}_i^0)$, and $\Delta(\Omega_i^0)$ are as defined in Section 3. In particular, $\Delta(\widetilde{\Omega}_i^0)$ and $\Delta(\Omega_i^0)$ are the sets of probability measures on the measurable space $(\widetilde{\Omega}_i^0, \mathscr{B}(\widetilde{\Omega}_i^0))$ and $(\Omega_i^0, \mathscr{B}(\Omega_i^0))$, respectively, where the Borel σ-algebras $\mathscr{B}(\widetilde{\Omega}_i^0)$ and $\mathscr{B}(\Omega_i^0)$ are as defined in Section 3. An analogous remark pertains to the higher-order spaces discussed below.

Since the topologies defined here distinguish belief hierarchies only on the basis of their beliefs about players' depth of reasoning (and not about the state of nature or players' signals), these topologies are weaker than the corresponding ones considered in Section 3. This means, of course, that the σ-algebra $\mathscr{B}_{H_i}$ on H_i generated by the auxiliary topology τ_{H_i} is (strictly) coarser than the usual Borel σ-algebra $\mathscr{B}(H_i)$ generated by the standard product topology τ_i on H_i (cf. Section 3).

The following result will be helpful:

Lemma A.1. The belief map ψ_i from the topological space (H_i, τ_{H_i}) to the topological space $(\Delta(S \times H_{-i}), \tau_{\Delta(S \times H_{-i})})$ is continuous.

Proof. Fix a belief hierarchy $h_i \in H_i$ and a net $\{h_i^\lambda\}_{\lambda \in \Lambda}$ such that $\{h_i^\lambda\}_{\lambda \in \Lambda}$ converges to h_i in τ_{H_i}. We want to show that $\{\psi_i(h_i^\lambda)\}_{\lambda \in \Lambda}$ converges to $\psi_i(h_i)$ in $\tau_{\Delta(S \times H_{-i})}$. We give the proof for the case that $h_i \in H_i^\infty$; the proof for the case that $h_i \in H_i^k$ for $k < \infty$ is similar. Since $h_i^\lambda \to h_i$, there is $\lambda' \in \Lambda$ such that $h_i^\lambda \in H_i^\infty$ for all $\lambda \geq \lambda'$. For $\lambda \geq \lambda'$, write $h_i^\lambda = (\mu_i^{\lambda,0}, \mu_i^{\lambda,1}, \mu_i^{\lambda,2}, \ldots)$; also, write $(\mu_i^0, \mu_i^1, \mu_i^2, \ldots)$ for h_i. Then, we have that for all $k = 0, 1, \ldots$, the sequence $\{\mu_i^{\lambda,k}\}_\lambda$ converges to μ_i^k in $\tau_{\Delta(\widetilde{\Omega}_i^k)}$. Since ψ_i is canonical, this is equivalent to $\{\mathrm{marg}_{\widetilde{\Omega}_i^{k-1}} \psi_i(h_i^\lambda)\}_\lambda$ converging to $\psi_i(h_i)$. The result now follows since the cylinders are a convergence-determining class. $\square$

As we have seen, a similar result holds for the standard topology (Corollary 3.3).

To emphasize, we use these auxiliary topologies purely as a tool to define the space $\mathcal{T}^\varepsilon$ (whose features we use in an intermediate step in the proof of Theorem 5.4). In particular, we do not change the topology on the space H_i of belief hierarchies (which is the standard topology τ_i defined in Section 3[24]), unlike the literature on strategic topologies (e.g., Dekel et al., 2006). Hence, all our results pertain to the standard (product) topology τ_i on H_i, $i = 1, 2$.

In the next section, we will use the topologies defined here to formalize the condition that players' beliefs about the other's depth of reasoning be independent of its belief about the other's signal.

A.2 Construction

Let $\varepsilon > 0$. To construct the space $\mathcal{T}^\varepsilon$, we need some more notation. For $i = 1, 2$ and $x_i \in X_i$, let $H_i(x_i; \varepsilon)$ be the set of types $h_i \in H_i$ with signal x_i such that their belief about the state of nature and the other player's signal is as defined in (1) above. Let $H_i(\varepsilon) :=$

---

[24]To wit, the topologies in Section 3 are defined as outlined in Section 2, with product spaces being endowed with the product topology, disjoint unions of spaces being endowed with the sum topology, and the space of Borel probability measures being endowed with the weak convergence topology.

$\bigcup_{x_i \in X_i}(\{(x_i, h_i^{*,0})\} \cup H_i(x_i; \varepsilon))$. Finally, for $h_i \in H_i$, let $\chi_i^*(h_i) \in X_i$ be the signal associated with the belief hierarchy $(x_i, \mu_i^0, \ldots) \in H_{-i}$

We construct $\mathcal{T}^\varepsilon$ from a family of belief-closed subspaces of the universal space $\mathcal{T}^*$. Let $(H_i')_{i=1,2}$ be a belief-closed subspace of $\mathcal{T}^*$ such that for each player $i = 1, 2$ and each type $h_i \in H_i' \setminus H_i^0$, there is $x_i \in X_i$ such that

  (i) beliefs about the state of nature and the other's signal are obtained from the signal x_i, that is, $h_i \in H_i(x_i; \varepsilon)$;

and if the depth $d_i(h_i)$ of h_i is at least 2, then

  (ii) The support of the belief $\psi_i(h_i)$ of h_i lies in $S \times H_{-i}(\varepsilon)$;

  (iii) h_i believes that the other player's beliefs are obtained from a signal x_{-i} if and only if the other player's signal is indeed x_{-i}, that is, for each signal $x_{-i} \in X_{-i}$,

  $$\psi_i(h_i)(x_{-i}, H_{-i}(x_{-i}; \varepsilon)) = \psi_i(h_i)(x_{-i}) = \psi_i(H_{-i}(x_{-i}; \varepsilon)).$$

  (iv) beliefs about players' depth of reasoning are independent of beliefs about signals, i.e., for $Y_{-i} \in \mathscr{B}(X_{-i})$ and $F_{-i} \in \mathscr{B}_{H_{-i}}$, we have that

  $$\psi_i(h_i)\big(F_{-i} \cap \{h_{-i} \in H_{-i} : \chi_{-i}^*(h_{-i}) \in Y_{-i}\}\big) = $$
  $$\psi_i(h_i)\big(F_{-i}\big)\, \psi_i(h_i)\big(\{h_{-i} \in H_{-i} : \chi_{-i}^*(h_{-i}) \in Y_{-i}\}\big). \quad (A.1)$$

For $i = 1, 2$, let T_i^ε be the union of the sets H_i' that belong to such a belief-closed subspace $(H_j')_{j=1,2}$. This defines the belief-closed subspace $\mathcal{T}^\varepsilon$. By the results in Appendix II, this belief-closed subspace defines a type space for players with a finite or infinite depth.

Note that the signal of a type $h_i = (x_i, \mu_i^0, \ldots) \in T_i^\varepsilon \cap H_i(x_i; \varepsilon)$ is $\chi_i^*(h_i) = x_i$. For $x_i \in X_i$, denote the set of types $T_i^\varepsilon \cap H_i(x_i; \varepsilon)$ that have received signal x_i by $T_i^\varepsilon(x_i)$.

Hence, the beliefs of the types in $\mathcal{T}^\varepsilon$ are consistent with the information structure, and, in fact, the information structure is common belief. Moreover, the belief of each type in T_i^ε about the other player's signal is independent of its belief about players' depth of reasoning. By construction, $\mathcal{T}^\varepsilon$ is the maximal belief-closed subspace which satisfies these properties.[25]

---

[25]It follows from the results in the online appendix (Appendix II) that $\mathcal{T}^\varepsilon$ is in fact a type space (for players with a finite or infinite depth of reasoning), and that any type space that also satisfies these properties can be mapped into $\mathcal{T}^\varepsilon$ via a unique belief-preserving morphism.

Appendix B   Equivalent beliefs about players' depth

It will be useful to define a collection $\mathcal{C}_i$, $i = 1, 2$, of classes of belief hierarchies for player i that are characterized by their beliefs about players' depth of reasoning.[26] For $i = 1, 2$, let $\mathcal{C}_i$ be the partition of H_i that contains the set $X_i \times \{h_i^{*,0}\}$ and the subsets

$$\left\{ h_i \in H_i : d_i(h_i) = d_i(h_i'), \psi_i(h_i)(E) = \psi_i(h_i')(E) \text{ for all } E \in \mathscr{B}_{H_{-i}} \right\} : \qquad h_i' \in H_i$$

of types that have the same depth of reasoning and the same beliefs about players' depth of reasoning. That is, every element C_i of $\mathcal{C}_i$ is an *(equivalence) class* of types with the same depth of reasoning and the same beliefs about players' depth of reasoning.

Let $\pi_i : H_i \to \mathcal{C}_i$ be the function that maps each $h_i \in H_i$ into the (unique) partition element of $\mathcal{C}_i$ to which it belongs. Let $\tau_{\mathcal{C}_i}$ be the strongest topology on $\mathcal{C}_i$ that makes the function π_i continuous with respect to the topology τ_{H_i} on H_i, that is,

$$\tau_{\mathcal{C}_i} := \left\{ F \subseteq \mathcal{C}_i : (\pi_i)^{-1}(F) \in \tau_{H_i} \right\}.$$

Since the function π_i is continuous with respect to τ_{H_i} (given the topology $\tau_{\mathcal{C}_i}$ on $\mathcal{C}_i$), it is also continuous with respect to the finer topology τ_i. We denote the Borel σ-algebra associated with $\tau_{\mathcal{C}_i}$ by $\mathscr{B}(\mathcal{C}_i)$.

We claim that the topological space $(\mathcal{C}_i, \tau_{\mathcal{C}_i})$ is homeomorphic to the universal space $\mathcal{T}^d$ (where d stands for "depth") with nontrivial beliefs only about players' depth of reasoning (and trivial beliefs about the state of nature and about signals). That is, let S^d be some arbitrary singleton $\{s^d\}$, and for $i = 1, 2$, let X_i^d be an arbitrary singleton $\{x_i^d\}$. Then, we can construct a universal type space on the set S^d of states of nature and signal spaces X_1^d and X_2^d as in the online appendix (Appendix II), giving a set H_i^d of belief hierarchies for each player $i = 1, 2$, endowed with the usual topology τ_i^d (defined analogously to the topology τ_i in Section 3). Then, it is straightforward to show by induction that the space $\mathcal{C}_i$, endowed with the topology $\tau_{\mathcal{C}_i}$, is homeomorphic to the topological space (H_i^d, τ_i^d), for $i = 1, 2$. It follows that $\mathcal{C}_i$ is a Polish space under the topology $\tau_{\mathcal{C}_i}$.

Finally, it will be useful to define a probability measure associated with each equivalence class. For $i = 1, 2$ and $C_i \in \mathcal{C}_i$ such that $C_i \neq X_i \times \{h_i^{*,0}\}$, let $\Psi_i(C_i)$ be the probability measure on $\mathscr{B}(\mathcal{C}_i)$ defined by:

$$\Psi_i(C_i)(E) := \psi_i(h_i)\left(\{ h_{-i} \in H_{-i} : \pi_{-i}(h_{-i}) \in E \} \right)$$

for $E \in \mathscr{B}(\mathcal{C}_i)$ and some arbitrary $h_i \in C_i$. Since π_{-i} is measurable and since the types in C_i have the same beliefs on $\mathscr{B}_{H_{-i}}$, this belief is well-defined. As $\mathcal{C}_{-i}$ is Polish when it is endowed with the topology $\tau_{\mathcal{C}_{-i}}$, the probability measure $\Psi_i(C_i)$ is regular.

---

[26]While we construct the classes for the case of two players, the construction can easily be generalized to the case of an arbitrary finite set of players.

Appendix C Proofs

C.1 Proof of Theorem 5.2 (continued)

We show that generic games with rich payoffs are responsive. That is, the set of responsive games is open and dense in the set of all two players, two actions games with rich payoffs (and continuous payoff functions, for fixed S). Denote the set of games with rich payoffs by $\mathcal{R} \subseteq \Gamma$, and recall that $\mathcal{R}$ is endowed with the relative topology induced by the sup-norm topology on Γ. We write $\mathcal{R}_i$ for the collection of utility functions u_i for i such that $(u_i, u'_{-i}) \in \mathcal{R}$ for some utility function u'_{-i} for the other player. The set of responsive games is denoted by $\mathcal{Y} \subseteq \mathcal{R}$, and $\mathcal{Y}_i$, $i = 1, 2$, is the set of utility functions $u_i \in \mathcal{R}_i$ for i such that $(u_i, u'_{-i}) \in \mathcal{Y}$ for some utility function $u'_{-i} \in \mathcal{R}_{-i}$ for the other player.

For $i = 1, 2$, we write $A_i = \{a_i, b_i\}$ for the action set for player i. For $s^{a_i}, s^{b_i} \in S$ such that $s^{a_i} \neq s^{b_i}$, let $\mathcal{R}_i(s^{a_i}, s^{b_i})$ be the set of utility functions in $\mathcal{R}_i$ such that a_i is strictly dominant at s^{a_i} and b_i is strictly dominant at s^{b_i}; such states exist by the rich payoffs assumption.

Fix $i = 1, 2$, $a_i \in A_i$, $s^{a_i}, s^{b_i} \in S$ such that $s^{a_i} \neq s^{b_i}$, distinct actions $a_{-i}, b_{-i} \in A_{-i}$, and $p_i \in [0, 1]$. Define $\mathcal{Y}_i(a_i, a_{-i}, b_{-i}, s^{a_i}, s^{b_i}, p_i) \subseteq \mathcal{R}_i(s^{a_i}, s^{b_i})$ to be the set of utility functions for player i such that a_i is strictly dominant at s^{a_i}, b_i is strictly dominant at s^{b_i}, and conditions (5.1) and (5.2) hold for this combination of actions and the probability p_i.

We claim that the collection $\mathcal{Y}_i(a_i, a_{-i}, b_{-i}, s^{a_i}, s^{b_i}, p_i)$ is open in $\mathcal{R}_i$. It then follows that the collection of responsive games is open, since this set is the finite product of unions of the sets $\mathcal{Y}_i(a_i, a_{-i}, b_{-i}, s^{a_i}, s^{b_i}, p_i)$, $a_i \in A_i$, $a_{-i}, b_{-i} \in A_{-i}$, $s^{a_i}, s^{b_i} \in S$, and $p_i \in [0, 1]$.

It suffices to show that for each $u_i \in \mathcal{Y}_i(a_i, a_{-i}, b_{-i}, s^{a_i}, s^{b_i}, p_i)$, there is an open set in $\mathcal{Y}_i(a_i, a_{-i}, b_{-i}, s^{a_i}, s^{b_i}, p_i)$ that contains u_i. Fix $u_i \in \mathcal{Y}_i(a_i, a_{-i}, b_{-i}, s^{a_i}, s^{b_i}, p_i)$. Then there exists $\delta > 0$ sufficiently small such that any $u'_i \in \mathcal{R}_i$ with $u'_i(c_i, d_{-i}, s^{e_i}) \in (u_i(c_i, d_{-i}, s^{e_i}) - \delta, u_i(c_i, d_{-i}, s^{e_i}) + \delta)$, $c_i \in A_i, d_{-i} \in A_{-i}$ and $s^{e_i} = s^{a_i}, s^{b_i}$, satisfies conditions (5.1) and (5.2). Consequently, $\mathcal{Y}_i(a_i, a_{-i}, b_{-i}, s^{a_i}, s^{b_i}, p_i)$ is open.

We next show that the set of responsive games is dense in $\mathcal{R}$. Again, it suffices to consider $\mathcal{R}_i$, $i = 1, 2$. Let $i = 1, 2$ and $u_i \in \mathcal{R}_i$. Let V be an open neighborhood of u_i. We want to show that the intersection of V with the set $\mathcal{Y}_i$ of "responsive" payoff functions for i is nonempty. If $u_i \in \mathcal{Y}_i$, then this is immediate; so suppose $u_i \in \mathcal{R}_i \setminus \mathcal{Y}_i$. Without loss of generality, assume that $u_i \in \mathcal{R}_i(s^{a_i}, s^{b_i})$. Since V is open, there is $\delta > 0$ sufficiently small such that any $u'_i \in \mathcal{R}_i$ with $u'_i(c_i, d_{-i}, s^{e_i}) \in (u_i(c_i, d_{-i}, s^{e_i}) - \delta, u_i(c_i, d_{-i}, s^{e_i}) + \delta)$, $c_i \in A_i, d_{-i} \in A_{-i}$ and $s^{e_i} = s^{a_i}, s^{b_i}$, belongs to V, and the intersection of the set of such utility functions u'_i and $\mathcal{Y}_i$ is nonempty. $\square$

C.2 Proof of Theorem 5.4

Recall that the usual (product) topology on H_i, $i = 1, 2$, is denoted by τ_i (Section 3). Also recall the definition of the auxiliary topology τ_{H_i}, $i = 1, 2$, from Appendix A. It is easy to see that the topology τ_{H_i} is coarser than τ_i (as the former only takes into account beliefs about depth of reasoning, and the latter considers all beliefs). Finally, recall the definition in Appendix B of the partition $\mathcal{C}_i$ of H_i into equivalence classes of types with the same depth of reasoning and the same (higher-order) beliefs about players' depth of reasoning; recall that $\mathcal{C}_i$ is endowed with the topology $\tau_{\mathcal{C}_i}$.

To prove the result, we need to consider the behavior of types with a finite depth. As in the proof of Theorem 5.2, it will be useful to introduce some extra notation. For $i = 1, 2$ and $\alpha = 0, 1, \ldots$, let $M_i^\alpha := H_i^\alpha$ be the set of types of depth α. To accommodate the sets B_i^n, $n = 0, 1, \ldots$, we need to consider the infinite ordinals $\omega, \omega + 1, \ldots$.[27] Let $M_i^\omega := B_i^0$, and for $n = 1, 2, \ldots$, let $M_i^{\omega+n} := B_i^n$. Thus, we consider the sequence of sets $\{M_i^\alpha\}_{\alpha < \omega + \omega}$.

The next lemma shows that for each player $i = 1, 2$ and $\alpha < \omega + \omega$, the set M_i^α is measurable and that there is $C_i^\alpha \in \mathcal{C}_i$ such that $C_i^\alpha \subseteq M_i^\alpha$. This implies that the sets M_i^α are well-defined, and that the statement in Theorem 5.4 is not void.

Lemma C.1. For every player $i = 1, 2$ and ordinal $\alpha < \omega + \omega$, we have that $M_i^\alpha \in \mathscr{B}_{H_i}$ (and thus $M_i^\alpha \in \mathscr{B}(H_i)$). Moreover, the set M_i^α is a (nonempty) union of elements of $\mathcal{C}_i$.

Proof. For $\alpha = 0$, M_i^0 is an element of $\mathcal{C}_i$ by definition; moreover, $M_i^0 \in \mathscr{B}_{H_i}$. For $\alpha > 0$, suppose that for all $\gamma < \alpha$, we have that $M_i^\gamma \in \mathscr{B}_{H_i}$ and that M_i^γ is a union of elements of $\mathcal{C}$. If α is finite, then it is immediate that M_i^α is a union of elements of $\mathcal{C}_i$ and that $M_i^\alpha \in \mathscr{B}_{H_i}$. So suppose $\alpha \geq \omega$. Then there is $G_{-i} \in \mathscr{B}_{H_{-i}}$ such that

$$M_i^\alpha = \left\{ h_i \in H_i^\infty : \psi_i(h_i)(G_{-i}) = 1 \right\}.$$

It is immediate that M_i^α is a union of elements of $\mathcal{C}_i$. That $M_i^\alpha \in \mathscr{B}_{H_i}$ follows from the the fact that $H_i^\infty \in \mathscr{B}_{H_i}$ and from Lemma A.1. Thus, we have $M_i^\alpha \in \mathscr{B}_{H_i}$ for all $\alpha < \omega + \omega$; since $\mathscr{B}(H_i)$ is finer that $\mathscr{B}_{H_i}$, it follows that $M_i^\alpha \in \mathscr{B}(H_i)$ for all $\alpha < \omega + \omega$. $\square$

In fact, this result implies that we could have strengthened the statement of Theorem 5.4 to apply to classes of beliefs that are merely *consistent* with nth-order mutual belief in infinite depth (i.e., class $C_i \in \mathcal{C}_i$ is such that $C_i \cap B_i^n \neq \emptyset$), as opposed to classes of beliefs that *satisfy* nth-order mutual belief in infinite depth (i.e., $C_i \subseteq B_i^n$): by Lemma C.1, the two are equivalent.

---

[27] Recall that $\omega = \{0, 1, \ldots\}$ is the first infinite ordinal, and that $\omega + n$, $n = 1, 2, \ldots$, is the successor of $\omega + n - 1$. The ordinal $\omega + \omega$ is the first ordinal that is greater than $0, 1, \ldots, \omega, \omega + 1, \omega + 2, \ldots$.

Recall that for every $\varepsilon > 0$, player $i = 1, 2$, and signal $x_i \in X_i$, $T_i^\varepsilon(x_i)$ is the set of types in T_i^ε that have received signal x_i. We have the following observation:

Lemma C.2. For every $\varepsilon > 0$, player $i = 1, 2$, signal $x_i \in X_i$, and class $C_i \in \mathcal{C}_i$, there is $h_i \in C_i$ such that $h_i \in T_i^\varepsilon(x_i)$.

In words, for every signal, and every "kind" of belief about players' depth of reasoning, as given by an equivalence class in $\mathcal{C}_i$, there is a type in $\mathcal{T}^\varepsilon$ with that signal and this "kind" of beliefs about players' depth of reasoning. The proof of Lemma C.2 follows directly from the definitions. We will use this result without mention.

The proof of Theorem 5.4 relies on the following lemma.

Lemma C.3. For every player $i = 1, 2$, ordinal $1 \le \alpha < \omega + \omega$, $y_i \in (0, \frac{1}{2})$, $C_i^\alpha \in \mathcal{C}_i$ such that $C_i^\alpha \subseteq M_i^\alpha$, there is $\varepsilon^\alpha(y_i, C_i^\alpha) > 0$ and an open neighborhood $O_i^\alpha(y_i, C_i^\alpha)$ of C_i^α in $\mathcal{C}_i$ such that for every $\varepsilon \le \varepsilon^\alpha(y_i, C_i^\alpha)$, both actions are strictly rationalizable for every type $h_i \in T_i^\varepsilon(x_i)$ with $x_i \in [y_i, \frac{1}{2}]$ and $\pi_i(h_i) \in O_i^\alpha(y_i, C_i^\alpha)$.

Proof. Recall that for each player $i = 1, 2$, both actions are rationalizable for the types in $M_i^0 = X_i \times \{h_i^{*,0}\}$.

Let $\alpha = 1$, $i = 1, 2$, $y_i \in (0, \frac{1}{2})$, and $C_i^1 \in \mathcal{C}_i$ such that $C_i^1 \subseteq M_i^1$. Then, $\Psi_i(C_i^1)(\{M_{-i}^0\}) = 1$. Define

$$O_i^1(y_i, C_i^1) := \{C_i \in \mathcal{C}_i : \Psi_i(C_i)(\{M_{-i}^0\}) > \Psi_i(C_i)(\{M_{-i}^0\}) - y_i\}.$$

Clearly, $C_i^1 \in O_i^1(y_i, C_i^1)$, so that $O_i^1(y_i, C_i^1)$ is nonempty. Using that $\mathcal{C}_i$ is homeomorphic to the set of belief hierarchies H_i^d (endowed with the usual topology) when there is uncertainty only about players' depth of reasoning (Appendix B), we obtain that the set $O_i^1(y_i, C_i^1)$ is open in $\tau_{\mathcal{C}_i}$ (Billingsley, 1968, App. III).

Let $\varepsilon > 0$ and $x_i \in [y_i, \frac{1}{2}]$. Let $h_i \in T_i^\varepsilon(x_i)$ be such that $\pi_i(h_i) \in O_i^1(y_i, C_i^1)$. Then,

$$\psi_i(h_i)(M_{-i}^0) > 1 - y_i,$$

and under the conjecture that all types in M_{-i}^0 invest (regardless of the state of nature), the expected payoff of investing to h_i is strictly greater than $(1 - y_i)x_i + y_i(x_i - 1) \ge 0$. Under the conjecture that the types in M_{-i}^0 do not invest, the expected payoff of investing to h_i is strictly less than $y_i x_i + (1 - y_i)(x_i - 1) = x_i + y_i - 1 < 0$. Hence, both actions are strictly rationalizable for h_i.

For $\alpha > 1$, suppose the claim is true for all $\gamma < \alpha$. We prove the result for the case that α is finite or $\alpha = \omega$; the proof for the case that $\alpha = \omega + n$, $n = 1, 2, \ldots$, is similar and therefore omitted.

Let $i = 1, 2$, $y_i \in (0, \frac{1}{2})$, and $C_i^\alpha \in \mathcal{C}_i$ such that $C_i^\alpha \subseteq M_i^\alpha$. Let $\Psi_i(C_i^\alpha)$ be the probability measure defined on $\mathscr{B}_{H_{-i}}$ defined as in Appendix B. Since $\Psi_i(C_i^\alpha)$ is regular, there is a subset

$$K_{-i}^{<\alpha}(y_i) \subseteq \{C_{-i} \in \mathcal{C}_{-i} : C_{-i} \subseteq M_{-i}^\gamma \text{ for some } \gamma < \alpha\}$$

that is compact (in $\tau_{\mathcal{C}_{-i}}$) and is such that

$$\Psi_i(C_i^\alpha)(K_{-i}^{<\alpha}(y_i)) > 1 - \frac{y_i}{3}.$$

By the induction hypothesis, for each $C_{-i}^\gamma \in K_{-i}^{<\alpha}(y_i)$, there is $\varepsilon^\gamma(\frac{y_i}{2}, C_{-i}^\gamma) > 0$ and an open neighborhood $O_{-i}^\gamma(\frac{y_i}{2}, C_{-i}^\gamma)$ of C_{-i}^γ in $\mathcal{C}_{-i}$ such that for all $\varepsilon \leq \varepsilon^\gamma(\frac{y_i}{2}, C_{-i}^\gamma)$, both actions are (strictly) rationalizable for the types in $h_{-i} \in T_{-i}^\varepsilon(x_{-i})$ with $x_{-i} \in [\frac{y_i}{2}, \frac{1}{2}]$ and $\pi_{-i}(h_{-i}) \in O_{-i}^\gamma(\frac{y_i}{2}, C_{-i}^\gamma)$.

Since $K_{-i}^{<\alpha}(y_i)$ is compact, there are finitely many $C_{-i}^{\gamma_1}, \ldots, C_{-i}^{\gamma_m} \in K_{-i}^{<\alpha}(y_i)$ such that

$$K_{-i}^{<\alpha}(y_i) \subseteq \bigcup_{\ell=1}^{m} O_{-i}^{\gamma_\ell}(\tfrac{y_i}{2}, C_{-i}^{\gamma_\ell}) := V_{-i}^{<\alpha}(y_i).$$

Note that $V_{-i}^{<\alpha}(y_i)$ is open in $\tau_{\mathcal{C}_{-i}}$ (and thus belongs to $\mathscr{B}(\mathcal{C}_{-i})$). Let

$$\tilde{\varepsilon}^\alpha(y_i, C_i^\alpha) := \min\{\varepsilon^{\gamma_1}(\tfrac{y_i}{2}, C_{-i}^{\gamma_1}), \ldots, \varepsilon^{\gamma_1}(\tfrac{y_i}{2}, C_{-i}^{\gamma_m})\},$$

so $\tilde{\varepsilon}^\alpha(y_i, C_i^\alpha) > 0$. Also, let

$$O_i^\alpha(y_i, C_i^\alpha) := \{C_i \in \mathcal{C}_i : \Psi_i(C_i)(V_{-i}^{<\alpha}(y_i)) > \Psi_i(C_i^\alpha)(V_{-i}^{<\alpha}(y_i)) - \tfrac{y_i}{3}\}.$$

It is immediate that $C_i^\alpha \in O_i^\alpha(y_i, C_i^\alpha)$, so the set is nonempty. Again, the set $O_i^\alpha(y_i, C_i^\alpha)$ is open in $\tau_{\mathcal{C}_i}$.

Define

$$Q_{-i}^{<\alpha}(y_i) := \{h_{-i} \in H_{-i} : \chi_i^*(h_{-i}) \in [\tfrac{y_i}{2}, \tfrac{1}{2}], \pi_{-i} \in V_{-i}^{<\alpha}(y_i)\}$$

to be the set of types in class $V_{-i}^{<\alpha}(y_i)$ that have received a signal in $[\frac{y_i}{2}, \frac{1}{2}]$; note that $Q_{-i}^{<\alpha}(y_i)$ is nonempty and that $Q_{-i}^{<\alpha}(y_i) \in \mathscr{B}(H_{-i})$. Then, for $\varepsilon > 0$ and $h_i \in T_i^\varepsilon(x_i)$ with $x_i \in [y_i, \frac{1}{2}]$ and $\pi_i(h_i) \in O_i^\alpha(y_i, C_i^\alpha)$, we have that

$$
\begin{aligned}
\psi_i(h_i)\big(Q_{-i}^{<\alpha}(y_i)\big) &= \psi_i(h_i)\big(\{h_{-i} \in H_{-i} : \chi_{-i}^*(h_{-i}) \in [\tfrac{y_i}{2}, \tfrac{1}{2}]\}\big)\, \psi_i(h_i)\big((\pi_{-i})^{-1}(V_{-i}^{<\alpha}(y_i))\big) \\
&> \psi_i(h_i)\big(\{h_{-i} \in H_{-i} : \chi_{-i}^*(h_{-i}) \in [\tfrac{y_i}{2}, \tfrac{1}{2}]\}\big)\big(1 - \tfrac{y_i}{3} - \tfrac{y_i}{3}\big),
\end{aligned}
$$

where the equality uses that beliefs about players' depth of reasoning are independent of beliefs about signals (condition (A.1) above). The inequality uses the definition of $V_{-i}^\alpha(y_i)$. By choosing ε small enough, the probability that the signal x_{-i} of a type $h_{-i} = (x_{-i}, \mu_i^0, \ldots) \in T_{-i}^\varepsilon$

33

is in the interval $[\frac{y_i}{2}, \frac{1}{2}]$ is arbitrarily close to 1. Consequently, there is positive $\varepsilon^\alpha(y_i, C_i^\alpha) \le \tilde{\varepsilon}^\alpha(y_i, C_i^\alpha)$ such that for all $\varepsilon \le \tilde{\varepsilon}^\alpha(y_i, C_i^\alpha)$

$$\psi_i(h_i)\big(Q_{-i}^{<\alpha}(y_i)\big) > 1 - y_i.$$

Let $\varepsilon \le \varepsilon^\alpha(y_i, C_i^\alpha)$, and let $h_i \in T_i^\varepsilon(x_i)$ be such that $x_i \in [y_i, \frac{1}{2}]$ and $\pi_i(h_i) \in O_i^\alpha(y_i, C_i^\alpha)$. Then, under the conjecture that all types in $Q_{-i}^{<\alpha}(y_i; \varepsilon)$ invest (for any state of nature), the expected payoff of investing to h_i is strictly greater than $(1 - y_i)x_i + y_i(x_i - 1) \ge 0$. Under the conjecture that the types in $Q_{-i}^{<\alpha}(y_i; \varepsilon)$ do not invest, the expected payoff of investing to (x_i, h_i) is strictly less than $y_i x_i + (1 - y_i)(x_i - 1) = x_i + y_i - 1 < 0$. Hence, both actions are strictly rationalizable for h_i. $\qquad\square$

By a symmetric argument, we can prove the "mirror image" of Lemma C.3: For every player $i = 1, 2$, ordinal $1 \le \alpha < \omega + \omega$, $y_i \in (\frac{1}{2}, 1)$, $C_i^\alpha \in \mathcal{C}_i$ such that $C_i^\alpha \subseteq M_i^\alpha$, there is $\varepsilon^\alpha(y_i, C_i^\alpha) > 0$ and an open neighborhood $O_i^\alpha(y_i, C_i^\alpha)$ of C_i^α in $\mathcal{C}_i$ such that for every $\varepsilon \le \varepsilon^\alpha(y_i, C_i^\alpha)$, both actions are strictly rationalizable for every type $h_i \in T_i^\varepsilon(x_i)$ with $x_i \in [\frac{1}{2}, y_i]$ and $\pi_i(h_i) \in O_i^\alpha(y_i, K_{-i}^{<\alpha}(y_i)_i^\alpha)$.

We are now ready to prove Theorem 5.4. If $\alpha = 0$, then $M_i^\alpha = X_i \times \{h_i^{*,0}\} = H_i^0$ for $i = 1, 2$; moreover, if $C_i^0 \subseteq M_i^0$, then $C_i^0 = M_i^0$. Since both actions are rationalizable for the types in H_i^0, we have that for any player $i = 1, 2$, signal $x_i \in (0, 1)$, $\varepsilon > 0$, and every type $h_i \in T_i^\varepsilon(x_i)$ with beliefs in class C_i^0, both actions are rationalizable for h_i, and the same is true for the types in the set

$$O_i^0(h_i) := H_i^0.$$

Clearly, $O_i^0(x_i, h_i)$ is an open neighborhood of h_i (in τ_i), and the result follows for $\alpha = 0$.

So suppose $0 < \alpha < \omega + \omega$. Fix a player $i = 1, 2$, a signal $x_i \in (0, 1)$, and a class $C_i^\alpha \in \mathcal{C}_i$ such that $C_i^\alpha \subseteq M_i^\alpha$. First suppose $x_i \le \frac{1}{2}$. By Lemma C.3, there is $\varepsilon^\alpha(x_i, C_i^\alpha) > 0$ and an open neighborhood $O_i^\alpha(x_i, C_i^\alpha)$ of C_i^α in $\mathcal{C}_i$ such that for every $\varepsilon \le \varepsilon^\alpha(x_i, C_i^\alpha)$, both actions are strictly rationalizable for every type $h_i \in T_i^\varepsilon(x_i)$ with $\pi_i(h_i) \in O_i^\alpha(x_i, C_i^\alpha)$.

Let $\varepsilon \le \varepsilon^\alpha(x_i, C_i^\alpha)$ be positive and let $h_i \in T_i^\varepsilon(x_i)$ such that $\pi_i(h_i) \in C_i^\alpha$. Since both actions are strictly rationalizable for h_i, there is $\delta > 0$ such that both actions are rationalizable for the types in

$$O_i^\alpha(h_i) := \Big\{h_i' \in H_i : \int_S s\, d\psi_i(h_i') \in (x_i - \delta, x_i + \delta), \pi_i(h_i') \in O_i^\alpha(x_i, C_i^\alpha)\Big\}.$$

Clearly, this set contains h_i, and it is open in τ_i (Kechris, 1995, e.g.,). For $x_i > \frac{1}{2}$, the result follows from a symmetric argument. $\qquad\square$

References

Angeletos, G.-M., C. Hellwig, and A. Pavan (2006). Signaling in a global game: Coordination and policy traps. *Journal of Political Economy 114*, 452–484.

Angeletos, G.-M., C. Hellwig, and A. Pavan (2007). Dynamic global games of regime change: Learning, multiplicity, and the timing of attacks. *Econometrica 75*, 711–756.

Angeletos, G.-M. and I. Werning (2006). Crises and prices: Information aggregation, multiplicity, and volatility. *American Economic Review 96*, 1720–1736.

Baliga, S. and T. Sjöström (2004). Arms races and negotiations. *Review of Economic Studies 71*, 351–369.

Battigalli, P., A. Di Tillio, E. Grillo, and A. Penta (2011). Interactive epistemology and solution concepts for games with asymmetric information. *The B.E. Journal of Theoretical Economics 11*(1 (Advances)), Article 6.

Battigalli, P. and M. Siniscalchi (2003). Rationalization and incomplete information. *Advances in Theoretical Economics 3*.

Bergemann, D. and S. Morris (2005). Robust mechanism design. *Econometrica 73*, 1771–1813.

Billingsley, P. (1968). *Convergence of Probability Measures*. New York: Wiley.

Brandenburger, A. and E. Dekel (1993). Hierarchies of beliefs and common knowledge. *Journal of Economic Theory 59*, 189–198.

Cabrales, A., R. Nagel, and R. Armenter (2007). Equilibrium selection through incomplete information in coordination games: An experimental study. *Experimental Economics 10*, 221–234.

Camerer, C. F., T.-H. Ho, and J.-K. Chong (2004). A cognitive hierarchy model of games. *Quarterly Journal of Economics 119*(3), 861–898.

Carlsson, H. and E. van Damme (1993). Global games and equilibrium selection. *Econometrica 61*, 989–1018.

Chen, Y.-C., A. Di Tillio, E. Faingold, and S. Xiong (2010). Uniform topologies on types. *Theoretical Economics 5*, 445–478.

Chen, Y.-C., S. Takahashi, and S. Xiong (2013). The robust selection of rationalizability. Working paper, National University of Singapore and University of Bristol.

Cole, H. L. and T. J. Kehoe (2000). Self-fulfilling debt crises. *Review of Economic Studies 67*, 91–116.

Corsetti, G., A. Dasgupta, S. Morris, and H. S. Shin (2004). Does one Soros make a difference? A theory of currency crises with large and small traders. *Review of Economic Studies 71*, 87–113.

Costa-Gomes, M., V. P. Crawford, and B. Broseta (2001). Cognition and behavior in normal-form games: An experimental study. *Econometrica 69*, 1193–1235.

Crawford, V. P., M. A. Costa-Gomes, and N. Iriberri (2012). Structural models of nonequilibrium strategic thinking: Theory, evidence, and applications. *Journal of Economic Literature*. Forthcoming.

Crawford, V. P. and N. Iriberri (2007). Level-k auctions: Can a nonequilibrium model of strategic thinking explain the winners curse and overbidding in private-value auctions? *Econometrica 75*, 1721–1770.

Dekel, E., D. Fudenberg, and S. Morris (2006). Topologies on types. *Theoretical Economics 1*, 275–309.

Dekel, E., D. Fudenberg, and S. Morris (2007). Interim correlated rationalizability. *Theoretical Economics 2*, 15–40.

Diamond, D. and P. Dybvig (1983). Bank runs, deposit insurance, and liquidity. *Journal of Political Economy 91*, 401–419.

Goldstein, I. and A. Pauzner (2005). Demand-deposit contracts and the probability of bank runs. *Journal of Finance 60*, 1293–1327.

Harsanyi, J. (1967–1968). Games on incomplete information played by Bayesian players. Parts I, II, III. *Management Science 14*, 159–182, 320–334, 486–502.

Heifetz, A. (1993). The Bayesian formulation of incomplete information–The non-compact case. *International Journal of Game Theory 21*, 329–338.

Heifetz, A. and W. Kets (2011). All types naive and canny. Working paper.

Heifetz, A. and D. Samet (1998). Topology-free typology of beliefs. *Journal of Economic Theory 82*, 324–341.

Heinemann, F., R. Nagel, and P. Ockenfels (2004). The theory of global games on test: Experimental analysis of coordination games with public and private information. *Econometrica 72*, 1583–1599.

Heinemann, F., R. Nagel, and P. Ockenfels (2009). Measuring strategic uncertainty in coordination games. *Review of Economic Studies 76*, 181–221.

Ho, T.-H., C. Camerer, and K. Weigelt (1998). Iterated dominance and iterated best response in experimental "*p*-beauty contests". *American Economic Review 88*, 947–969.

Hocking, J. G. and G. S. Young (1988). *Topology.* Mineola, NY: Dover Publications.

Kechris, A. S. (1995). *Classical Descriptive Set Theory.* Graduate Texts in Mathematics. Berlin: Springer-Verlag.

Kets, W. (2010). Bounded reasoning and higher-order uncertainty. Working paper, Northwestern University.

Kets, W. (2013). Finite depth of reasoning and equilibrium play in games with incomplete information. Working paper, Northwestern University.

Mertens, J. F. and S. Zamir (1985). Formulation of Bayesian analysis for games with incomplete information. *International Journal of Game Theory 14*, 1–29.

Monderer, D. and D. Samet (1989). Approximating common knowledge with common beliefs. *Games and Economic Behavior 1*, 170–190.

Morris, S. and H. S. Shin (1998). Unique equilibrium in a model of self-fulfilling currency attacks. *American Economic Review 88*, 587–597.

Morris, S. and H. S. Shin (2003). Global games: Theory and applications. In M. Dewatripont, L. P. Hansen, and S. J. Turnovsky (Eds.), *Advances in economics and econometrics: Eighth World Congress*, Volume 1. Cambridge University Press.

Morris, S. and H. S. Shin (2004). Liquidity black holes. *Review of Finance 8*, 1–18.

Nagel, R. (1995). Unraveling in guessing games: An experimental study. *American Economic Review 85*, 1313–1326.

Obstfeld, M. (1986). Rational and self-fulfilling balance-of-payments crises. *American Economic Review 76*, 72–81.

Obstfeld, M. (1996). Models of currency crises with self-fulfilling features. *European Economic Review 40*, 1037–1047.

Parthasarathy, K. R. (1978). *Introduction to probability and measure.* New York: Springer-Verlag.

Purves, R. (1966). Bimeasurable functions. *Fundamenta Mathematicae 58*, 149–157.

Rubinstein, A. (1989). The electronic mail game: Strategic behavior under "almost common knowledge". *American Economic Review 79*, 385–391.

Stahl, D. O. and P. W. Wilson (1995). On players' models of other players: Theory and experimental evidence. *Games and Economic Behavior 10*, 218–254.

Strzalecki, T. (2009). Depth of reasoning and higher order beliefs. Working paper, Harvard University.

Takeuti, G. (1987). *Proof Theory*, Volume 81 of *Studies in Logic and the Foundations of Mathematics.* Elsevier.

Weinstein, J. and M. Yildiz (2007). A structure theorem for rationalizability with application to robust predictions of refinements. *Econometrica 75*, 365–400.

Yang, M. (2013). Coordination with flexible information acquisition. Working paper, Duke University.