THE STRATEGIC IMPACT OF HIGHER-ORDER BELIEFS

By

Yi-Chun Chen, Alfredo Di Tillio,
Eduardo Faingold, and Siyang Xiong

September 2012

COWLES FOUNDATION DISCUSSION PAPER NO. 1875

COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281

http://cowles.econ.yale.edu/

The Strategic Impact of Higher-Order Beliefs*

Yi-Chun Chen† Alfredo Di Tillio‡ Eduardo Faingold§ Siyang Xiong¶

August 23, 2012

Abstract

Previous research has established that the predictions made by game theory about strategic behavior in incomplete information games are quite sensitive to the assumptions made about the players' infinite hierarchies of beliefs. We evaluate the severity of this robustness problem by characterizing conditions on the primitives of the model—the players' hierarchies of beliefs—for the strategic behavior of a given Harsanyi type to be approximated by the strategic behavior of (a sequence of) perturbed types. This amounts to providing characterizations of the strategic topologies of Dekel, Fudenberg, and Morris (2006) in terms of beliefs. We apply our characterizations to a variety of questions concerning robustness to perturbations of higher-order beliefs, including genericity of common priors, and the connections between robustness of strategic behavior and the notion of common p-belief of Monderer and Samet (1989).

Keywords: Games with incomplete information, rationalizability, higher-order beliefs, robustness

JEL Classification: C70, C72.

*We thank Pierpaolo Battigalli, Eddie Dekel, Jeffrey Ely, Barton Lipman, George Mailath, Hervé Moulin, Drew Fudenberg, Atsushi Kajii, Stephen Morris, Daisuke Oyama, Marcin Pęski, Muhamet Yildiz, and participants at various seminars and conferences for helpful comments. Eduardo Faingold and Siyang Xiong thank the National Science Foundation (grants SES-1227593 and SES-1227620) for financial support.

†Dept. of Economics, National University of Singapore, ecsycc@nus.edu.sg

‡Dept. of Economics and IGIER, Università Luigi Bocconi, alfredo.ditillio@unibocconi.it

§Dept. of Economics and Cowles Foundation, Yale University, eduardo.faingold@yale.edu

¶Dept. of Economics, Rice University, xiong@rice.edu

1 Introduction

A major concern with non-cooperative game theory is its reliance on details. The formal description of a strategic situation as a game requires minutiae that are often not available to the analyst in real life, such as the precise order of moves, the information and actions available to the players when they move, and the exact payoff functions and beliefs of the players. Unfortunately, game theoretic solutions are known to depend very sensitively on those details, and the search for robust predictions—those that are approximately correct across a range of similar models—becomes a necessity.

Concentrating on strategic-form games with incomplete information (Harsanyi, 1967-68), we investigate the robustness of game theoretic predictions to misspecification of the players' subjective beliefs. We take the point of view of the analyst who posits a Harsanyi type space to model the players' uncertainty and recognizes that his model is only an idealization, thus being subject to misspecification error. For example, the analyst may assume that there is common knowledge of the underlying state of the world, but understands that such common knowledge assumption can be at best an approximation of reality. Or, as is often the case in practice, the analyst may posit a non-degenerate type space with a common prior, but realizes that the true common prior distribution may be slightly different from the one assumed, or even that the players may have slightly different priors.

What is the impact of such kinds of misspecifications on the analyst's behavioral predictions? We attack this question by studying the tail properties of the *hierarchies of beliefs* encoded in Harsanyi types (the beliefs of a player about the payoff-relevant parameters, his beliefs about the other players' beliefs about the payoff-relevant parameters, and so on, ad infinitum) and their implications for behavior. Our main finding is an exact characterization, in terms of the primitives of the model (the players' hierarchies of beliefs), of what it takes for a pair of types to display similar strategic behaviors. Thus we find the minimum level of precision of the analyst's information model that is required for accurate predictions of strategic play.

To explain our results we first need to be precise about what we mean by "strategic behavior." Our behavioral assumption is that players play a Bayesian equilibrium on a type space (possibly without a common prior). Thus, from the perspective of the analyst, who does not know the true type space of the players and has a concern for robustness, the relevant solution concept is (interim correlated) rationalizability (Dekel, Fudenberg, and Morris, 2006). Indeed, the set of actions that are rationalizable for a type t coincides with the set of actions that can be played in some Bayesian equilibrium on some type space, by some type that has the same hierarchy of beliefs as t (Dekel, Fudenberg, and Morris, 2007, Remark 2). A similar perspective is taken by Bergemann and Morris (2009) in the context of robust mechanism design. See also Aumann (1987) and Brandenburger and Dekel (1987) for early papers pioneering this approach.

Formally, our main results are characterizations of the *strategic topology* and the *uniform* *strategic topology* of Dekel, Fudenberg, and Morris (2006). The former is the coarsest topology on the universal type space (Mertens and Zamir, 1985)—the space of all hierarchies of beliefs—under which the correspondence that maps each type of a player into his set of rationalizable actions displays the same kind of continuity properties that the best-reply, Nash equilibrium and rationalizability correspondences exhibit in complete information games.[1] Thus, for any player, a sequence of types t^n converges in the strategic topology to a type t if and only if, for every finite game and every action a of the player in the game, the following conditions are equivalent: (a) action a is rationalizable for type t; (b) for every $\varepsilon > 0$ and sufficiently large n, action a is ε-rationalizable for type t^n, where ε is a size of sub-optimization allowed in the incentive constraints. Convergence in the uniform strategic topology adds the requirement that the rate of convergence in (b) be uniform across all finite games (with uniformly bounded payoffs).

As shown by Dekel, Fudenberg, and Morris (2006), a sequence of types converges in the strategic topology only if it converges in the *product topology*: for every integer $k \geqslant 1$, the sequence of k-order beliefs must converge weakly. However, the Electronic Mail game of Rubinstein (1989) and, more generally, the structure theorem of Weinstein and Yildiz (2007), show that convergence in the product topology does not imply strategic convergence. Our characterizations are based on a strengthening of product convergence that requires k-order beliefs to converge at a rate that is uniform in k.

We first explain the characterization of the uniform strategic topology, as it is simpler to state and can serve as a benchmark for the other characterization result. For each k, endow the space of k-order beliefs with the Prohorov distance, which is a standard distance that metrizes the topology of weak convergence of probability measures (Billingsley, 1999). Say that a sequence of types t^n converges *uniform-weakly* to a type t if the k-order belief of t^n converges to the k-order belief of t and the rate of convergence is uniform in k. Our first main result, Theorem 1, states that uniform strategic convergence is equivalent to uniform weak convergence.[2] To interpret, consider an analyst who would like to make predictions with some minimal level of accuracy, and wants to achieve this level of accuracy *uniformly* across all strategic situations that the players might face.[3] A tight condition for such uniformly robust prediction is that the analyst's model of the players' beliefs and higher-order beliefs be sufficiently precise, with the required degree of precision, *as measured by the Prohorov distance*, binding uniformly over all levels of the belief hierarchy.

The content of Theorem 1 can be dissected in two parts. First, the theorem underscores the role of uniform convergence of hierarchies of beliefs as a requirement for robustness. In light of the structure theorem of Weinstein and Yildiz (2007), which shows that the tails of the hierarchies

[1] See the introduction of Dekel, Fudenberg, and Morris (2006) for a precise analogy with complete information games.

[2] The partial result that uniform weak convergence implies uniform strategic convergence was proved in Chen, Di Tillio, Faingold, and Xiong (2010). The reverse implication is new to the present paper.

[3] This may be the case if the analyst is a mechanism designer who will ultimately determine the game that the players will face by his choice of a mechanism.

of beliefs can have a large impact on strategic behavior, the role of uniform convergence should not come as a surprise. Second, the theorem *quantifies* the impact of a misspecification at each order of the hierarchy by the Prohorov distance. We view this as a nontrivial part of the theorem. Indeed, the Prohorov distance, on which the notion of uniform weak convergence is based, is but one of many equivalent distances that metrize the topology of weak convergence of probability measures. For any such distance, one can consider the associated uniform distance over infinite hierarchies of beliefs. It turns out that these distances may generate different topologies over infinite hierarchies, even though the induced topologies over k-order beliefs coincide for each finite k.[4] Theorem 1 identifies one of these uniform distances that ultimately characterizes the uniform-strategic topology.

The characterization of the strategic topology (Theorem 2) is also based on uniform convergence and the Prohorov metric, but is more subtle. The relevant class of events for uniform weak convergence, and a fortiori, uniform strategic convergence, is the entire Borel σ-algebra of the universal type space. By contrast, our characterization of the strategic topology highlights the role of coarser information structures called *frames*. A frame is a profile of finite partitions of the universal type space—one partition for each player—that satisfies a measurability condition: each player's belief concerning the events in the frame must pin down a unique atom of that player's partition. (We discuss the meaning of this condition below.) For any frame $\mathcal{P}$ and any positive integer k, we define a distance over types, $d^k_{\mathcal{P}}$, that is analogous to the Prohorov distance over k-order beliefs, but restricts the events for which the proximity is measured to those in the frame $\mathcal{P}$. Say that a sequence of types t^n converges to a type t *uniform-weakly on* $\mathcal{P}$ if, for every positive integer k, t^n converges to t under $d^k_{\mathcal{P}}$ and the rate of convergence is uniform in k. Our second main result, Theorem 2, states that a sequence of types converges strategically if and only if it converges uniform-weakly on every frame.

A frame can be interpreted as a coarsening of the canonical information structure of the universal type space where each player is assumed to know only his belief about the payoff-relevant states and the events in the other players' partitions. The measurability condition is a natural requirement for a model of coarse information in a multi-agent setting: it amounts to the condition that all interactive knowledge events (i knows j knows $\ldots$) are measurable with respect to the players' partitions. In particular, in a frame, every event concerning player j that player i can reason about is either known to be true or known to be false by player j. Since the strategic behavior of a player in a game can only be affected by events that the player knows, it is intuitive that strategic convergence should imply uniform-weak convergence only on frames and not on all information structures.

To shed further light on the impact of higher-order beliefs, we use our characterization to

[4]In Chen, Di Tillio, Faingold, and Xiong (2010) we report an example of a sequence of types that converges uniform weakly but fails to converge in the uniform topology associated with a distance (different from Prohorov) that metrizes the topology of weak convergence of probability measures.

investigate the connection between strategic convergence and a natural notion of uniform convergence based on *common* p*-beliefs* (Monderer and Samet, 1989).[5] Say that a sequence of types t^n *converges in common beliefs* to a type t if t^n converges to t in the product topology and, for every event E and every $p \geqslant 0$, the following conditions are equivalent: (i) E is common p-belief for type t; (ii) for every $\varepsilon > 0$, $k \geqslant 1$ and sufficiently large n, type t^n has common $(p-\varepsilon)$-belief on the event that the players have k-order beliefs that are ε-close to those from E. This is the interim analogue of the ex ante notion of convergence based on common p-beliefs that the seminal papers of Monderer and Samet (1996) and Kajii and Morris (1998) have shown to characterize the ex ante strategic topology for Bayesian equilibrium on countable common prior type spaces.

We establish, using the characterization of Theorem 2, that strategic convergence implies convergence in common beliefs (Theorem 3). But, somewhat surprisingly, we find that the converse fails: we exhibit a sequence of types that converges in common beliefs but does not converge uniform-weakly on a frame (Example 5). These results highlight a fundamental difference between the common prior, equilibrium, ex ante framework of the early literature and our non-common prior, non-equilibrium, interim framework. Nonetheless, when the limit is assumed a *finite type*—a type that belongs to a finite type space—we show that convergence in common beliefs is equivalent to uniform weak convergence, and hence, a fortiori, to both uniform strategic and strategic convergence (Theorem 4).

Finally, we use our characterizations to revisit, and reverse, two important genericity results concerning the structure of the universal type space. The first result, due to Ely and Pęski (2011), shows that critical types—those types which display discontinuous rationalizable behavior—form a meager set under the product topology. By way of contrast, we show that they form an open and dense set under the strategic topology (Theorem 5). Second, Lipman (2003) shows that types consistent with a common prior are dense in the universal type space under the product topology. Instead, we show that those types are nowhere dense under the strategic topology (Theorem 6). We also report measure-theoretic versions of these genericity results based on the notions of prevalence and shyness.

The rest of the paper is organized as follows. Section 2 sets up the incomplete information model and reviews the basic definitions and properties of type spaces, hierarchies of beliefs, common p-beliefs and the solution concept of interim correlated rationalizability. Section 3 presents the strategic topologies of Dekel, Fudenberg, and Morris (2006) and their characterizations in terms of beliefs, and studies their relationship with the notion of common p-belief of Monderer and Samet (1989). Section 4 examines the genericity of critical types and common prior types under the strategic topology and also their measure-theoretic genericity. Section 5 concludes by discussing possible extensions for future work. All proofs are presented in the appendix.

[5] Recall that an event E is common p-belief for a given type if that type assigns probability at least p to E, assigns probability at least p to the event that E obtains and the other players assign probability at least p to E, and so forth, ad infinitum.

2 Preliminaries

Given a measurable space S, write $\Delta(S)$ for the space of probability measures on S, equipped with the σ-algebra generated by the sets of the form $\{\nu \in \Delta(S) : \nu(E) \geqslant p\}$, where $E \subseteq S$ is a measurable set and $p \in [0, 1]$. Unless otherwise stated, product spaces are assumed to be endowed with the product σ-algebra, subspaces with the relative σ-algebra, and finite spaces with the discrete σ-algebra.

We consider two-player games with incomplete information.[6] The space Θ of *payoff-relevant states* (or just *states*, for short) is assumed to be finite, with $\#\Theta \geqslant 2$, and is fixed throughout the paper. A *game* is a profile $G = (A_i, g_i)_{i \in I}$, where A_i is a finite set of actions of player $i \in I = \{1, 2\}$ and $g_i : A_i \times A_{-i} \times \Theta \to \mathbb{R}$ is his payoff function, which is extended to $\Delta(A_i) \times \Delta(A_{-i} \times \Theta)$ in the usual way.[7] The uncertainty of the players is modeled by a *type space*, that is, a profile $(T_i, \phi_i)_{i \in I}$ where T_i is a measurable space of types of player i and $\phi_i : T_i \to \Delta(\Theta \times T_{-i})$ is a measurable map that associates, with each type t_i of player i, his belief $\phi_i(t_i)$ about the payoff-relevant states and the types of player $-i$.

2.1 Solution Concept

Our analysis is based on the solution concept of *interim correlated rationalizability*, due to Dekel, Fudenberg, and Morris (2007). The concept is an incomplete information extension of the rationalizability of Bernheim (1984) and Pearce (1984).

Given a game $G = (A_i, g_i)_{i \in I}$ and $\varepsilon \geqslant 0$, an action $a_i \in A_i$ is an *ε-best reply* to $\pi_{-i} \in \Delta(A_{-i} \times \Theta)$, written $a_i \in BR_i(\pi_{-i}, G, \varepsilon)$, if $g_i(a_i, \pi_{-i}) \geqslant g_i(a'_i, \pi_{-i}) - \varepsilon$ for all $a'_i \in A_i$. Given a type space $(T_i, \phi_i)_{i \in I}$, a profile of correspondences $\varsigma_i : T_i \rightrightarrows A_i$, $i \in I$, has the *ε-best-reply property* if for each $i \in I$, $t_i \in T_i$ and $a_i \in \varsigma_i(t_i)$ there is a *conjecture* $\nu \in \Delta(\Theta \times T_{-i} \times A_{-i})$ such that the following conditions hold:

$$\operatorname{marg}_{\Theta \times T_{-i}} \nu = \phi_i(t_i),$$
$$\big(\operatorname{marg}_{T_{-i} \times A_{-i}} \nu\big)\big[\operatorname{graph} \varsigma_{-i}\big] = 1,\text{[8]}$$
$$a_i \in BR_i\big(\operatorname{marg}_{A_{-i} \times \Theta} \nu, G, \varepsilon\big).$$

The greatest (w.r.t. pointwise set inclusion) profile of correspondences that has the ε-best-reply property is the *interim correlated ε-rationalizable correspondence*—or *ε-rationalizable correspondence*, for short—and is denoted $R_i(\cdot, G, \varepsilon) : T_i \rightrightarrows A_i$.[9] For $a_i \in R_i(t_i, G, \varepsilon)$, we say

[6] We restrict attention to two-player games for ease of notation. All our results extend to the general N-player case.

[7] Following standard notation, for each player $i \in I$ we write $-i$ to designate the other player in I.

[8] When graph ς_i is not measurable, this expression is taken to mean that $\operatorname{marg}_{T_{-i} \times A_{-i}} \nu_i$ assigns probability one to a measurable subset of graph ς_i.

[9] Such greatest profile of correspondences is well defined, because the pointwise union of any family of profiles

that a_i is an *ε-rationalizable action for type t_i*. Finally, to ease notation, for $\varepsilon = 0$ we denote the correspondence $R_i(\cdot, G, 0)$ simply by $R_i(\cdot, G)$ and call it the *(interim correlated) rationalizable correspondence*.

2.2 Beliefs

Type spaces contain implicit descriptions of each player's beliefs about the payoff-relevant states, his beliefs about the other player's beliefs about the payoff-relevant states, and so on. We formulate this higher-order uncertainty following Mertens and Zamir (1985), whose construction we now review. An equivalent formulation is found in Brandenburger and Dekel (1993).

A *first-order belief* of player i is a probability distribution over Θ, that is, an element of $\mathcal{T}_i^1 = \Delta(\Theta)$. Recursively, for $k \geqslant 2$, the space of *k-order beliefs* of player i is

$$\mathcal{T}_i^k = \Big\{ \big(t_i^1, \dots, t_i^{k-1}, t_i^k\big) \in \mathcal{T}_i^{k-1} \times \Delta(\Theta \times \mathcal{T}_{-i}^{k-1}) \;:\; \operatorname{marg}_{\Theta \times \mathcal{T}_{-i}^{k-2}} t_i^k = t_i^{k-1} \Big\}.$$

The above condition on marginal distributions implies that each element of $\mathcal{T}_i^k$ is uniquely determined by its k-th coordinate, hence we can identify $\mathcal{T}_i^k$ with $\Delta(\Theta \times \mathcal{T}_{-i}^{k-1})$. The space of *hierarchies of beliefs* of player i is

$$\mathcal{T}_i = \big\{ \big(t_i^1, t_i^2, \dots\big) \;:\; (t_i^1, \dots, t_i^k) \in \mathcal{T}_i^k \quad \forall k \geqslant 1 \big\}.$$

Mertens and Zamir (1985) show that for every hierarchy $t_i = (t_i^1, t_i^2, \dots) \in \mathcal{T}_i$ there is a unique probability measure $\mu_i(t_i) \in \Delta(\Theta \times \mathcal{T}_{-i})$ that extends each of the measures $t_i^1, t_i^2, \dots$. Moreover, the map $\mu_i : \mathcal{T}_i \to \Delta(\Theta \times \mathcal{T}_{-i})$ is an isomorphism.[10] To ease notation, for each event $E \subseteq \Theta \times \mathcal{T}_{-i}$ we shall often write $\mu_i(E \,|\, t_i)$ instead of the more cumbersome $\mu_i(t_i)[E]$.

The type space $(\mathcal{T}_i, \mu_i)_{i \in I}$ is called the *universal type space*, since any type t_i from any abstract type space $(T_i, \phi_i)_{i \in I}$ uniquely induces a hierarchy of beliefs $\tau_i(t_i) = (\tau_i^1(t_i), \tau_i^2(t_i), \dots)$ in a natural way: $\tau_i^1(t_i) = \operatorname{marg}_\Theta \phi_i(t_i)$ and, recursively, for $k \geqslant 2$,

$$\tau_i^k(t_i)\big[\theta \times E\big] = \phi_i(t_i)\Big[\theta \times \big(\tau_{-i}^{k-1}\big)^{-1}(E)\Big]$$

for each $\theta \in \Theta$ and measurable $E \subseteq \mathcal{T}_{-i}^{k-1}$, as shown by Mertens and Zamir (1985).

We restrict attention throughout to types that belong to the universal type space $\mathcal{T}_i$. This incurs no loss of generality, as Dekel, Fudenberg, and Morris (2007) have shown that the ε-rationalizable

of correspondences with the ε-best-reply property must also have the ε-best-reply property. The ε-rationalizable correspondence has an alternative characterization in terms of iterated elimination of actions that are never an interim best-reply (Dekel, Fudenberg, and Morris, 2007, Claim 1) and a dual characterization in terms of iterated elimination of strongly interim-dominated actions (Chen, Di Tillio, Faingold, and Xiong, 2010, Proposition 1).

[10] That is, a measurable bijection with measurable inverse.

actions of any type are pinned down by its hierarchy of beliefs: $R_i(t_i, G, \varepsilon) = R_i\big(\tau_i(t_i), G, \varepsilon\big)$ for any type t_i belonging to any type space. Thus, we take $\mathcal{T}_i$ to be the domain of the correspondence $R_i(\cdot, G, \varepsilon)$.

2.3 Common Beliefs

A standard approach to express assumptions about higher-order uncertainty uses the notion of *common belief*, due to Monderer and Samet (1989, 1996). Let $\Omega = \Theta \times \mathcal{T}_1 \times \mathcal{T}_2$ and, for each measurable set $E \subseteq \Omega$ and each type $t_i \in \mathcal{T}_i$, let E_{t_i} designate the section of E over t_i, i.e.

$$E_{t_i} = \big\{(\theta, t_{-i}) \,:\, (\theta, t_1, t_2) \in E\big\},$$

which is a measurable set by standard arguments. For each $p \in [0, 1]$, define

$$B_i^p(E) = \big\{t_i \in \mathcal{T}_i \,:\, \mu_i(E_{t_i} \,|\, t_i) \geqslant p\big\},$$

which is also measurable.[11] Then, for each $\mathbf{p} = (p_1, p_2) \in [0, 1]^2$, define the event that E *is mutual* $\mathbf{p}$*-belief* as

$$B^{\mathbf{p}}(E) = \Theta \times B_1^{p_1}(E) \times B_2^{p_2}(E),$$

and the event that E *is common* $\mathbf{p}$*-belief* as

$$C^{\mathbf{p}}(E) \;=\; B^{\mathbf{p}}(E) \;\cap\; B^{\mathbf{p}}\big(E \cap B^{\mathbf{p}}(E)\big) \;\cap\; B^{\mathbf{p}}\Big(E \cap B^{\mathbf{p}}\big(E \cap B^{\mathbf{p}}(E)\big)\Big) \;\cap\; \cdots.$$

Then, the event that E *is common* $\mathbf{p}$*-belief for player* i, written $C_i^{\mathbf{p}}(E)$, is defined as the projection of $C^{\mathbf{p}}(E)$ on $\mathcal{T}_i$, which is a measurable set because $C^{\mathbf{p}}(E)$ is a rectangle.[12] For notational convenience, we identify $B_i^{p_i}(E)$ and $C_i^{\mathbf{p}}(E)$ with the cylinders $\Theta \times B_i^{p_i}(E) \times \mathcal{T}_{-i}$ and $\Theta \times C_i^{\mathbf{p}}(E) \times \mathcal{T}_{-i}$, respectively. Thus, we can write

$$C_i^{\mathbf{p}}(E) \;=\; B_i^{p_i}(E) \;\cap\; B_i^{p_i}\big(E \cap B_{-i}^{p_{-i}}(E)\big) \;\cap\; B_i^{p_i}\Big(E \cap B_{-i}^{p_{-i}}\big(E \cap B_i^{p_i}(E)\big)\Big) \;\cap\; \cdots.$$

Note that for $p_{-i} = 0$ we have $C_i^{\mathbf{p}}(E) = B_i^{p_i}(E)$.

Finally, common belief has the following well known fixed-point characterization.

Lemma 1. *$C^{\mathbf{p}}(E) = B^{\mathbf{p}}(E \cap C^{\mathbf{p}}(E))$ and $C_i^{\mathbf{p}}(E) = B_i^{p_i}(E \cap C_{-i}^{\mathbf{p}}(E))$.*

---

[11] The measurability of $B_i^p(E)$ follows from the measurability of the map $t_i \mapsto \mu_i(E_{t_i} \,|\, t_i)$. The class of events

$$\big\{E \subseteq \Omega \,:\, E \text{ is a measurable set such that } t_i \mapsto \mu_i(E_{t_i} \,|\, t_i) \text{ is measurable}\big\}$$

can be readily verified to be a monotone class containing the algebra of finite disjoint unions of measurable rectangles, which generates the product σ-algebra on Ω. It follows that the map $t_i \mapsto \mu_i(E_{t_i} \,|\, t_i)$ is measurable for every measurable E.

[12] The definition of $C^{\mathbf{p}}(E)$ is analogous to the *common repeated belief* of Monderer and Samet (1996), which differs from the original definition of Monderer and Samet (1989). A similar definition appears in Ely and Pęski (2011) for the case where E is a rectangle. We allow the event E to be any measurable set.

3 Behavior- and Belief-Based Topologies

Two kinds of information are encapsulated in the description of a type. On one hand, a type generates an infinite hierarchy of beliefs. On the other hand, a type determines the player's strategic behavior in each game. Accordingly, we distinguish between *belief-based* topologies on types—under which, by definition, nearby types have similar hierarchies of beliefs—and *behavior-based* topologies—under which nearby types display similar strategic behaviors. In this section, we provide belief-based characterizations of behavior-based topologies.

The *product topology* is the canonical belief-based topology on the universal type space $\mathcal{T}_i$. To define it, let the (finite-dimensional) space of first-order beliefs, $\mathcal{T}_i^1 = \Delta(\Theta)$, be endowed with the Euclidean topology and, recursively, for each $k \geqslant 2$, let the space of k-order beliefs, $\mathcal{T}_i^k = \Delta(\Theta \times \mathcal{T}_{-i}^{k-1})$, be endowed with the topology of weak convergence of probability measures relative to the topology defined on $\mathcal{T}_{-i}^{k-1}$. Then, a sequence of types t_i^n converges to a type t_i in the product topology if the k-order belief of t_i^n converges to the k-order belief of t_i, for every $k \geqslant 1$.[13]

Turning to behavior topologies, we first review the key definitions of Dekel, Fudenberg, and Morris (2006).

Definition 1 (Strategic topology)**.** A sequence of types t_i^n *converges strategically* to a type t_i if for every game G and every action a_i of player i in G, the following conditions are equivalent:

(a) a_i is rationalizable for t_i in G;

(b) for every $\varepsilon > 0$ there exists N such that for every $n \geqslant N$, a_i is ε-rationalizable for t_i^n in G.

The *strategic topology* is the topology of strategic convergence on $\mathcal{T}_i$.[14,15]

[13] Since Θ is finite, the Borel σ-algebra of the product topology on $\mathcal{T}_i$ coincides with the σ-algebra obtained in the topology-free formulation of Section 2.2. Likewise, the Borel σ-algebra on $\Delta(\Theta \times \mathcal{T}_{-i})$ generated by the topology of weak convergence of probability measures (induced by the product topology on $\mathcal{T}_{-i}$) coincides with the σ-algebra of our topology-free formulation. Finally, if we endow each $\mathcal{T}_i$ with the product topology and $\Delta(\Theta \times \mathcal{T}_{-i})$ with the topology of weak convergence of probability measures, then $\mathcal{T}_i$ and $\Delta(\Theta \times \mathcal{T}_{-i})$ are compact metrizable and $\mu_i : \mathcal{T}_i \to \Delta(\Theta \times \mathcal{T}_{-i})$ is a homeomorphism.

[14] That is, by definition, a set $F \subseteq \mathcal{T}_i$ is closed under the strategic topology if it contains the limit points of all strategically convergent sequences in F. Alternatively, the strategic topology can be defined as the topology generated by the collection of sets having either the form $\{t_i \in \mathcal{T}_i : a_i \notin R_i(t_i, G, \varepsilon)\}$ or the form $\{t_i \in \mathcal{T}_i : a_i \in \cup_{\varepsilon' < \varepsilon} R_i(t_i, G, \varepsilon')\}$ and can be shown to be metrizable.

[15] The definition of strategic convergence above follows Ely and Pęski (2011). The original definition of Dekel, Fudenberg, and Morris (2006) is different, but both definitions are equivalent. Under the original definition, a sequence $t_i^n \to t_i$ strategically if, for every game G, every action a_i in G and every $\gamma \geqslant 0$, the following are equivalent: (i) $a_i \in R_i(t_i, G, \gamma)$; (ii) $\forall \varepsilon > 0\ \exists N$ such that $\forall n \geqslant N$, $a_i \in R_i(t_i^n, G, \gamma + \varepsilon)$. The 2007 working paper version of Ely and Pęski (2011) proves the equivalence between the two definitions.

The implication from (b) to (a)—for every game G and action a_i of player i in G—is a form of upper hemi-convergence of the rationalizable correspondence, and is equivalent to convergence in the product topology (Dekel, Fudenberg, and Morris, 2006, Theorems 1 and 2). In turn, the implication from (a) to (b) is a form of lower hemi-convergence, which, instead, does *not* follow from convergence in the product topology, as the well known example of Rubinstein (1989) demonstrates. Finally, if the implication from (a) to (b) holds for every game G and action a_i of player i in G, then the implication from (b) to (a) also holds (Dekel, Fudenberg, and Morris, 2006, Corollary 1).

If we strengthen Definition 1 to require the rate of convergence N to be independent of the game G (in all games with uniformly bounded payoffs), we then obtain the following definition. Given a game $G = (A_i, g_i)_{i \in I}$, write $|G| = \max_i \max |g_i|$.

Definition 2 (Uniform strategic topology)**.** A sequence of types t_i^n *converges uniform-strategically* to a type t_i if for every payoff bound $M > 0$ there exist positive integers $(N_\varepsilon)_{\varepsilon>0}$ such that for every game G with $|G| \leqslant M$ and every action a_i of player i in G, the following conditions are equivalent:

(a) a_i is rationalizable for t_i in G;

(b) for every $\varepsilon > 0$ and $n \geqslant N_\varepsilon$, a_i is ε-rationalizable for t_i^n in G.

The *uniform strategic topology* is the topology of uniform strategic convergence on $\mathcal{T}_i$.

We turn to the characterizations of the strategic and the uniform-strategic topologies in terms of beliefs. We begin with the uniform-strategic topology, as its characterization takes a simpler form.

3.1 Characterization of the Uniform Strategic Topology

To characterize the uniform strategic topology, we use a notion of convergence of types under which the rate of convergence is uniform across the levels of the belief hierarchy. In order to define this uniformity, we first need to fix a distance on the space of k-order beliefs. We use the Prohorov distance, which metrizes the topology of weak convergence of probability measures.

For each integer $k \geqslant 1$, we define recursively a distance d_i^k on $\mathcal{T}_i$ as the Prohorov distance over k-order beliefs assuming that the space of $(k-1)$-order beliefs of player $-i$ is endowed with the distance d_{-i}^{k-1}. Thus, for each player i, we set $d_i^0 \equiv 0$ and, for each integer $k \geqslant 0$ and types s_i and t_i,

$$d_i^{k+1}(s_i, t_i) = \inf \left\{ \delta > 0 : \mu_i(E | \, t_i) \leqslant \mu_i \left(E^{\delta,k} \middle| \, s_i \right) + \delta \ \text{ for each measurable } E \subseteq \Theta \times \mathcal{T}_{-i} \right\},$$

where

$$E^{\delta,k} = \Big\{(\theta, s_{-i}) \in \Theta \times \mathcal{T}_{-i} \ : \ d^k_{-i}(s_{-i}, t_{-i}) < \delta \ \text{ for some } t_{-i} \text{ with } (\theta, t_{-i}) \in E\Big\}.$$[16]

We then consider the following notion of uniform convergence, introduced in Chen, Di Tillio, Faingold, and Xiong (2010).

Definition 3 (Uniform weak convergence)**.** A sequence of types t_i^n converges *uniform-weakly* to a type t_i if

$$d_i^{UW}(t_i^n, t_i) \stackrel{\text{def}}{=} \sup_{k\geqslant 1} d_i^k(t_i^n, t_i) \to 0 \quad \text{as} \quad n \to \infty.$$

Thus, uniform weak convergence is the uniform counterpart of product convergence when the topology of weak convergence of k-order beliefs is metrized by the Prohorov distance. In Chen, Di Tillio, Faingold, and Xiong (2010), we showed that uniform-weak convergence implies uniform-strategic convergence. Here we prove the reverse implication, thus establishing the equivalence:

Theorem 1. *A sequence of types converges uniform-strategically if and only if it converges uniform-weakly.*

The proof of this and all other results is presented in Appendix A.

Ever since Rubinstein's (1989) seminal paper, misspecifications of higher-order beliefs have been recognized to have a potentially large impact on strategic predictions. The systematic treatment of Weinstein and Yildiz (2007) exposed the pervasiveness of this sensitivity by showing that the phenomenon is not peculiar to the Electronic Mail game, and hence advocated wider scrutiny of the assumptions one makes about the players' subjective beliefs. Theorem 1 quantifies the exact impact of such assumptions (uniformly over games) by identifying the appropriate measure of proximity of hierarchies of beliefs. In effect, the role of the Prohorov distance in the definition of uniform weak convergence, and hence in our characterization result, turns out to be nontrivial. For any distance that metrizes the topology of weak convergence of probability measures, one can define an associated uniform distance over infinite hierarchies of beliefs. However, these distances may generate different topologies over infinite hierarchies, even though the induced topologies over k-order beliefs coincide for each finite k, as shown in an example in Chen, Di Tillio, Faingold, and Xiong (2010, Section 5.2). Theorem 1 identifies one of these uniform distances that ultimately characterizes the uniform-strategic topology.

Theorem 1 serves as a benchmark for our characterization of the strategic topology in the next section, which is also based on uniform convergence across the orders of the hierarchy of beliefs. Uniform weak convergence requires proximity of beliefs concerning *all* events in the universal type space. As will soon become apparent, relaxing this requirement is the key to our characterization of the strategic topology.

[16] Viewed as a distance on $\mathcal{T}_i$, d_i^k is only a pseudo-distance—as opposed to a standard distance—, since there exist distinct types with the same k-order beliefs (and hence different ℓ-order beliefs, for some $\ell > k$).

3.2 Characterization of the Strategic Topology

We begin with a definition that plays a central role in our characterization.

Definition 4 (Frames)**.** A *frame* is a profile $\mathcal{P} = (\mathcal{P}_i)_{i \in I}$, where each $\mathcal{P}_i$ is a finite measurable partition of $\mathcal{T}_i$, such that for each $s_i, t_i \in \mathcal{T}_i$,

$$\mu_i(\theta \times E \mid s_i) = \mu_i(\theta \times E \mid t_i) \quad \forall \theta \in \Theta,\ \forall E \in \mathcal{P}_{-i} \quad \Longrightarrow \quad s_i \in \mathcal{P}_i(t_i).^{17}$$

To interpret this condition, note that any profile of partitions induces a coarsening of the canonical information structure of the universal type space, as follows. Given a profile of partitions $\mathcal{P} = (\mathcal{P}_i)_{i \in I}$ (not necessarily a frame), say that a type t_i *knows* an event $E \subseteq \Theta \times \mathcal{T}_{-i}$ (relative to $\mathcal{P}$) if $\{s_i \in \mathcal{T}_i : \mu_i(\theta \times F|s_i) = \mu_i(\theta \times F|t_i)\ \forall \theta \in \Theta,\ \forall F \in \mathcal{P}_{-i}\} \subseteq E$. Thus, in this coarse information model, player i knows his beliefs about $\Theta \times \mathcal{P}_{-i}$ (positive introspection), and that is all he knows (hence the coarseness of the model). It is then readily verified that $\mathcal{P}$ is a frame if and only if every event concerning player $-i$ that player i can reason about is either known to be true or known to be false by player $-i$.[18] Equivalently, $\mathcal{P}$ is a frame if and only if all events of the form "i knows $-i$ knows $\ldots$ i knows E," where $E \in 2^{\Theta} \otimes \mathcal{P}_{-i}$, are measurable with respect to $\mathcal{P}_i$.

The notion of frame is key to our characterization, so it will be useful to go over a few examples to illustrate the definition. We begin by describing a general procedure for constructing a new frame from a given frame. We will then make use of the procedure to discuss three canonical examples of frames.

Given a frame $\mathcal{P}$ and a finite measurable partition Π_i of the finite-dimensional simplex $\Delta(\Theta \times \mathcal{P}_{-i})$, define the *partition on $\mathcal{T}_i$ induced by* Π_i, written $\mathcal{T}_i / \Pi_i$, as follows: any two types of player i belong to the same element of $\mathcal{T}_i / \Pi_i$ if and only if their beliefs over $\Theta \times \mathcal{P}_{-i}$ belong to the same element of Π_i. The following lemma is straightforward from the definitions:

Lemma 2. *Each $\mathcal{T}_i / \Pi_i$ is a measurable partition of $\mathcal{T}_i$, and the join $(\mathcal{P}_i \vee (\mathcal{T}_i / \Pi_i))_{i \in I}$ is a frame.*[19]

We now present the examples.

Example 1 (Finite-order frames)**.** A *finite-order frame* is a frame whose atoms are k-order measurable events, for some integer $k \geqslant 1$. For instance, any profile of first-order measurable partitions is a (first-order) frame, as it can be readily verified from Definition 4. Examples of higher

[17] Following standard notation, $\mathcal{P}_i(t_i)$ designates the atom of $\mathcal{P}_i$ containing t_i.

[18] That is, either player $-i$ knows the event, or he knows its complement.

[19] Recall that the join of a pair of partitions, denoted by the symbol $\vee$, is the coarsest partition that is finer than both partitions in the pair.

order frames can be constructed by successive application of Lemma 2 beginning with an arbitrary profile of first-order measurable partitions. For instance, given a $\theta_0 \in \Theta$ consider the following profile of first-order measurable bi-partitions:

$$\mathscr{P}_i = \Big\{ \{t_i : \mu_i(\theta_0|t_i) \in \mathbb{Q}\}, \{t_i : \mu_i(\theta_0|t_i) \notin \mathbb{Q}\} \Big\}, \quad i \in I,$$

where $\mathbb{Q}$ is the set of rational numbers. Then, consider the bi-partition Π_1 of the simplex $\Delta(\Theta \times \mathscr{P}_2)$ into the set

$$\Big\{ q \in \Delta(\Theta \times \mathscr{P}_2) : q\big(\theta_0, \{t_2 : \mu_2(\theta_0|t_2) \in \mathbb{Q}\}\big) \geqslant 1/2 \Big\}$$

and its complement. Thus, the join $\mathscr{P}_1 \vee (\mathcal{T}_1 / \Pi_1)$ partitions $\mathcal{T}_1$ into four second-order measurable events, according to whether or not a type t_1 satisfies each of the following two conditions:

$$\mu_1(\theta_0|t_1) \in \mathbb{Q}, \quad \mu_1\Big(\theta_0 \times \{t_2 : \mu_2(\theta_0|t_2) \in \mathbb{Q}\} \Big| t_1\Big) \geqslant 1/2.$$

The profile $(\mathscr{P}_1 \vee (\mathcal{T}_1/\Pi_1), \mathscr{P}_2)$ is an example of a second-order frame. ♦

Example 2 (Common belief frames)**.** Let $\mathbf{p} = (p_1, p_2) \in (0,1]^2$ and $E = \Theta' \times \mathcal{T}_1 \times \mathcal{T}_2$, where Θ' is a subset of Θ. The profile of bi-partitions $\mathscr{P}_i = \{C_i^{\mathbf{p}}(E), \mathcal{T}_i \setminus C_i^{\mathbf{p}}(E)\}$, $i \in I$, is a frame. Indeed, if any two types of player i agree on the probabilities over $\Theta \times \mathscr{P}_{-i}$, then they must agree on the probability of the event $\Theta' \times C_{-i}^{\mathbf{p}}(E) = E \cap C_{-i}^{\mathbf{p}}(E)$. Then, either both types assign probability at least p_i to $E \cap C_{-i}^{\mathbf{p}}(E)$, in which case both types belong to $C_i^{\mathbf{p}}(E)$ (by Lemma 1), or both types assign probability less than p_i to $E \cap C_{-i}^{\mathbf{p}}(E)$, in which case both types belong to the complement of $C_i^{\mathbf{p}}(E)$ (again, by Lemma 1).

More generally, given any frame $\mathscr{P}$ and any event $E \subseteq \Omega$ that is $\mathscr{P}$-measurable,[20] the join between $\mathscr{P}_i$ and $\{C_i^{\mathbf{p}}(E), \mathcal{T}_i \setminus C_i^{\mathbf{p}}(E)\}$ is a frame, called a *common belief frame*. ♦

Example 3 (Strategic frames)**.** Given $\varepsilon \geqslant 0$ and a game $G = (A_i, g_i)_{i \in I}$, the *ε-strategic frame* associated with G is the profile

$$\mathscr{P}_i = \{[B_i] : \varnothing \neq B_i \subseteq A_i\}, \quad i \in I,$$

where

$$[B_i] = \{t_i : R_i(t_i, G, \varepsilon) = B_i\}.$$

(For $\varepsilon = 0$ we call it simply the strategic frame of G.) To check that the above is indeed a frame, suppose that s_i and t_i are two arbitrary types of player i that have the same beliefs concerning all the events of the form $\theta \times [B_{-i}]$. Let us show that s_i and t_i have the same set of ε-rationalizable actions in G. If $a_i \in R_i(t_i, G, \varepsilon)$ then, by definition, there is a conjecture $\nu \in \Delta(\Theta \times \mathcal{T}_{-i} \times A_{-i})$ that

[20] This means that E is measurable with respect to the algebra $2^\Theta \otimes \mathscr{P}_1 \otimes \mathscr{P}_2$ on Ω.

satisfies $\mathrm{marg}_{\Theta\times\mathcal{T}_{-i}}\,\nu = \mu_i(t_i)$, $\nu\big(\operatorname{graph} R_{-i}(\cdot, G, \varepsilon)\big) = 1$ and $a_i \in BR_i\big(\mathrm{marg}_{A_{-i}\times\Theta}\,\nu, G, \varepsilon\big)$. Then, define a conjecture $\bar{\nu} \in \Delta(\Theta\times\mathcal{T}_{-i}\times A_{-i})$ for type s_i as follows: for each $\theta \in \Theta$, $a_{-i} \in A_{-i}$ and measurable $E \subseteq \mathcal{T}_{-i}$,

$$\bar{\nu}(\theta \times E \times a_{-i}) = \sum_{B_{-i}} \frac{\nu(\theta \times [B_{-i}] \times a_{-i})}{\nu(\theta \times [B_{-i}])}\, \mu_i(\theta \times (E \cap [B_{-i}])|s_i),$$

where the summation ranges over all $B_{-i} \subseteq A_{-i}$ such that $\mu_i(\theta\times(E\cap[B_{-i}])|s_i) > 0$. Notice that $\bar{\nu}$ is well defined, because whenever $\mu_i(\theta\times(E\cap[B_{-i}])|s_i) > 0$ we must also have $\nu(\theta\times[B_{-i}]) > 0$, since

$$\nu(\theta \times [B_{-i}]) = \mu_i(\theta \times [B_{-i}]|t_i) = \mu_i(\theta \times [B_{-i}]|s_i). \tag{1}$$

By construction, we have $\mathrm{marg}_{\Theta\times\mathcal{T}_{-i}}\,\bar{\nu} = \mu_i(s_i)$. Also, the condition $\nu\big(\operatorname{graph} R_{-i}(\cdot, G, \varepsilon)\big) = 1$ implies that for every $\theta \in \Theta$, $a_{-i} \in A_{-i}$ and $B_{-i} \subseteq A_{-i}$,

$$\bar{\nu}(\theta \times [B_{-i}] \times a_{-i}) > 0 \quad\Longrightarrow\quad \nu(\theta \times [B_{-i}] \times a_{-i}) > 0 \quad\Longrightarrow\quad a_{-i} \in B_{-i}.$$

Hence, $\bar{\nu}\big(\operatorname{graph} R_{-i}(\cdot, G, \varepsilon)\big) = 1$. Finally, (1) above implies $\mathrm{marg}_{A_{-i}\times\Theta}\,\bar{\nu} = \mathrm{marg}_{A_{-i}\times\Theta}\,\nu$. Thus, $a_i \in BR_i(\mathrm{marg}_{A_{-i}\times\Theta}\,\bar{\nu}, G, \varepsilon)$, and therefore $a_i \in R_i(s_i, G, \varepsilon)$. We have thus shown that $R_i(t_i, G, \varepsilon) \subseteq R_i(s_i, G, \varepsilon)$, and the opposite inclusion can be proved by interchanging the roles of s_i and t_i in the argument above. This proves that the ε-strategic frame is indeed a frame. ♦

We emphasize that *not every frame is the strategic frame of a game*. For instance, in a strategic frame, each player's partition must contain an atom that is open in the product topology, namely, any atom consisting of types whose set of rationalizable actions is minimal (w.r.t. set inclusion).[21] General frames, however, need not contain open sets. The first-order frame in Example 1 illustrates this fact.

Turning to the characterization of strategic convergence, we introduce a notion of uniform weak convergence of types *relative to a fixed frame* $\mathcal{P}$. For each player i, set $d^0_{i,\mathcal{P}} \equiv 0$ and, for each integer $k \geqslant 0$ and types s_i and t_i, define

$$d^{k+1}_{i,\mathcal{P}}(s_i, t_i) = \inf\Big\{\delta > 0 \ :\ \mu_i(E|\,t_i) \leqslant \mu_i\big(E^{\delta,k}_{\mathcal{P}}\,\big|\,s_i\big) + \delta \quad \forall E \in 2^{\Theta} \otimes \mathcal{P}_{-i}\Big\},$$

where

$$E^{\delta,k}_{\mathcal{P}} = \Big\{(\theta, s_{-i}) \in \Theta \times \mathcal{T}_{-i} \ :\ d^k_{-i,\mathcal{P}}(s_{-i}, t_{-i}) < \delta \ \text{ for some } t_{-i} \text{ with } (\theta, t_{-i}) \in E\Big\}.$$

Thus, the definition of $d^k_{i,\mathcal{P}}$ is similar to that of d^k_i, but restricts the events for which the proximity is measured to those in the frame $\mathcal{P}$.[22]

[21] That such an atom is product-open follows from minimality and the upper hemi-coninuity of the rationalizable correspondence.

[22] This restriction makes $d^k_{i,\mathcal{P}}$ only a pre-distance, that is, it satisfies $d^k_{i,\mathcal{P}}(s_i, t_i) \geqslant 0$ and $d^k_{i,\mathcal{P}}(t_i, t_i) = 0$, but it fails symmetry and the triangle inequality.

Definition 5 (Uniform weak convergence on frames)**.** A sequence of types t_i^n converges to a type t_i *uniform-weakly on a frame* $\mathcal{P}$ if

$$d_{i,\mathcal{P}}^{UW}(t_i^n, t_i) \stackrel{\text{def}}{=} \sup_{k \geqslant 1} d_{i,\mathcal{P}}^{k}(t_i^n, t_i) \to 0 \quad \text{as} \quad n \to \infty.$$

With this definition in place, we are ready to state our characterization.

Theorem 2. *A sequence of types converges strategically if and only if it converges uniform-weakly on every frame.*

The theorem has conceptual and practical significance. First, the result deepens our understanding of the belief underpinnings of strategic robustness by characterizing, in terms of the primitives of the model (hierarchies of beliefs), the class of all perturbations to which the predictions of rationalizability are robust.

Second, as in Theorem 1, the result draws attention to a *particular form* of uniform convergence across the levels of the belief hierarchy as a condition for robustness. While it is expected that some form of uniform convergence should play a role, it is much less clear at the outset what kind of uniformity would ultimately lead to a characterization. For instance, one might expect that the strategically relevant notion of uniform convergence were the one that requires every event that is common belief for the limit type to remain arbitrarily close to an event that is approximately common belief for the types that are sufficiently far in the tail of the sequence. Indeed, an ex ante variation of this condition characterizes the robust perturbations for Bayesian equilibrium in countable common prior type spaces, as shown in the early papers of Monderer and Samet (1996) and Kajii and Morris (1998). However, this is not the case in our framework, for that notion of convergence turns out to be equivalent to uniform weak convergence on common belief frames (c.f. Example 2), whereas the characterization above requires convergence on *all* frames, and the latter turns out to be a stronger condition.

Third, the theorem highlights the role of frames as the coarse information structures that are relevant for strategic convergence. The role of frames in the "only if" direction is intuitive. Given a profile of partitions that is not a frame, there is always a player i who can reason about some event E that concerns player $-i$ such that player $-i$ cannot know whether or not E obtains.[23] But since the strategic behavior of a player can only depend on events that the player knows, it is intuitive that failure of uniform weak convergence on a profile of partitions that is not a frame need not imply failure of strategic convergence. As for the "if" direction, the role of frames is more mechanical. If a sequence t_i^n converges uniform weakly on all frames, then, given an arbitrary game, the sequence must converge uniform-weakly on the strategic frame generated by that game. This fact, combined with the continuity of the ε-best-reply correspondence in the Prohorov metric,

[23] The precise meanings of "reason about" and "know", relative to the coarse information model induced by a profile of partitions, is explained in the paragraph that follows Definition 4.

leads naturally to an induction argument (over the levels of the hierarchy) that proves the strategic convergence of t_i, as in the proof presented in Appendix A.3.

Fourth, the theorem leads to a connection between the strategic topology and the early literature on robustness mentioned above, and in particular to the notion of common p-beliefs, which plays a prominent role in that literature. In the next section, we use Theorem 2 to show that strategic convergence implies an interim analogue of the mode of convergence based on common p-beliefs of the early literature. We also present a striking example that shows that the converse fails.

Fifth, the characterization enables the genericity analysis that we carry out in Section 4, where we revisit, and reverse, two important genericity results of the recent literature (due to Ely and Pęski (2011) and Lipman (2003)), which were carried out in the product topology.

Finally, there are instances when the characterization has practical significance. To illustrate, we present Example 4 below, where we describe a case in which it is particularly simple to apply the characterization to prove that a sequence converges strategically, even though the sequence does not converge uniform weakly. Nonetheless, we do not maintain that this is a general feature of the characterization, that is, we do not claim that our characterization has any computational advantages over the definition of strategic convergence in general. In terms of computational complexity, checking uniform weak convergence in every frame appears to be as hard as checking convergence of rationalizable behavior in every game. The main point of our exercise remains to express conditions for robustness in terms of primitives.

Example 4 (A sequence that converges uniform-weakly on all frames)**.** To build our example, we need a preliminary definition. Fix two distinct states $\theta_0,\ \theta_1 \in \Theta$ and define the *Dirac type space* (based on $\{\theta_0, \theta_1\}$) as the abstract type space $(X_i, \phi_i)_{i \in I}$, where each $X_i = \{0,1\}^{\mathbb{N}}$ (the space of all infinite sequences of 0's and 1's) and for each type $x_i = (x_{i1}, x_{i2}, \ldots)$, the belief $\phi_i(x_i)$ assigns probability one to the singleton $\{(\theta_{x_{i1}}, \ell_{-i}(x_i))\}$, where ℓ_{-i} is the left-shift operator: $\ell_{-i}\left(x_{i1}, x_{i2}, \ldots\right) = (x_{i2}, x_{i3}, \ldots)$. Thus, the hierarchy of beliefs of a Dirac type $x_i = (x_{i1}, x_{i2}, \ldots)$ assigns probability one to state $\theta_{x_{i1}}$, probability one to $-i$ assigning probability one to $\theta_{x_{i2}}$, probability one to $-i$ assigning probability one to i assigning probability one to $\theta_{x_{i3}}$, and so on.

We now present an example of a sequence of Dirac types, s_i^n, for which it will be particularly simple to prove uniform weak convergence on all frames, even though s_i^n will not converge uniform weakly. For each positive integer n, let b_n be the finite sequence of length $2n$ comprising n zeroes followed by n ones. Thus, $b_1 = (0,1)$, $b_2 = (0,0,1,1)$, and so on. Let t_i^n and s_i^n be the Dirac types of player i such that $t_i^n = (b_n, b_{n+1}, \ldots)$ and $s_i^n = (b_1, \ldots, b_{n-1}, 0, 0, 0, \ldots)$. It can be readily verified that $d_i^k(s_i^n, t_i^1)$ equals zero for $k \leqslant n^2$ and equals one for $k > n^2$. Therefore, the sequence s_i^n converges to t_i^1 in the product topology, but not uniform-weakly.

Let us show that s_i^n converges to t_i^1 uniform weakly on every frame. First, note that $\ell_{-i}(t_i^n)$

and $\ell_{-i}(s_i^1) = (0, 0, 0, \ldots)$ share the same first $n-1$ coordinates (a sequence of $n-1$ zeroes), hence $d_{-i}^{n-1}(\ell_{-i}(t_i^n), \ell_{-i}(s_i^1)) = 0$. It follows that the sequence $\ell_{-i}(t_i^n)$ converges to $\ell_{-i}(s_i^1)$ in the product topology. Next, fix an arbitrary frame $\mathcal{P}$. Let $\mathcal{P}'_{-i} \subseteq \mathcal{P}_{-i}$ be the set of all elements of $\mathcal{P}_{-i}$ that contain infinitely many types from the sequence $\ell_{-i}(t_i^n)$. Thus, for every n and $P_{-i} \in \mathcal{P}'_{-i}$, there is a type in P_{-i} whose first $n-1$ coordinates are all zeroes. It follows that, for every k and $P_{-i} \in \mathcal{P}'_{-i}$, we have $\ell_{-i}(s_i^1) \in (P_{-i})_{\mathcal{P}}^{0,k}$ and, therefore, $\mu_i(\theta_0 \times (P_{-i})_{\mathcal{P}}^{0,k} | s_i^1) = 1$. Furthermore, there exists N such that, for every $n \geqslant N$, there exists $P_{-i} \in \mathcal{P}'_{-i}$ such that $\ell_{-i}(t_i^n) \in P_{-i}$ and, therefore, $\mu_i(\theta_0 \times P_{-i} | t_i^n) = 1$. It follows that $d_{i,\mathcal{P}}^k(s_i^1, t_i^n) = 0$ for every $k \geqslant 1$ and $n \geqslant N$. Finally, since $s_i^n = (b_1, \ldots, b_{n-1}, s_i^1)$ and $t_i^1 = (b_1, \ldots, b_{n-1}, t_i^n)$ have the same first $2n(n+1)$ coordinates, a straightforward induction argument shows that $d_{i,\mathcal{P}}^{k+2n(n-1)}(s_i^n, t_i^1) = 0$ for every $k \geqslant 1$ and $n \geqslant N$. Thus, for all $n \geqslant N$, we have $\sup_{k \geqslant 1} d_{i,\mathcal{P}}^k(s_i^n, t_i^1) = 0$, and hence the sequence s_i^n converges to t_i^1 uniform-weakly on $\mathcal{P}$. As this is true for every frame $\mathcal{P}$, we conclude that s_i^n converges strategically to t_i^1, as claimed. ♦

3.3 Convergence in Common Beliefs

Previous work on strategic topologies for Bayesian equilibrium (BE) under common priors (Monderer and Samet, 1996; Kajii and Morris, 1998) has proved equivalences between an ex ante notion of strategic convergence for BE and ex ante notions of convergence based on common beliefs. In this section, we consider the interim version of common belief convergence, prove that it is a necessary condition for convergence in the strategic topology for interim correlated rationalizability, and demonstrate, by means of an example, that it fails to be sufficient. This highlights a fundamental difference between the strategic topologies in the common prior, equilibrium, ex ante framework of Monderer and Samet (1996) and Kajii and Morris (1998) and in our non-common prior, non-equilibrium, interim framework.

Definition 6 (Common belief convergence)**.** A sequence of types t_i^n *converges in common beliefs* to a type t_i if for every $\delta > 0$, every integer $k \geqslant 1$, every measurable set $E \subseteq \Omega$ and every $\mathbf{p} \in [0,1]^2$ with $t_i \in C_i^{\mathbf{p}}(E)$, there exists N such that for every $n \geqslant N$, $t_i^n \in C_i^{\mathbf{p}-\delta\mathbf{1}}(E^{\delta,k})$.[24]

To shed light on the definition, an analogy with product convergence is useful. A necessary and sufficient condition for a sequence of types t_i^n to converge to a type t_i in the product topology is that for every $\delta > 0$, every integer $k \geqslant 1$, every measurable set $E \subseteq \Omega$ and every $p \in [0,1]^2$ with $t_i \in B_i^p(E)$, there exists N such that for every $n \geqslant N$, $t_i^n \in B_i^{p-\delta\mathbf{1}}(E^{\delta,k})$.[25] There is thus

[24]If we modify the definition to require the equivalence between (a) and (b) to hold only for events E that are closed in the product topology, then the notion of convergence remains the same. Indeed, letting $\overline{E}$ denote the product-topology closure of E, we have $C_i^{\mathbf{p}}(E) \subseteq C_i^{\mathbf{p}}(\overline{E})$ and $E^{\delta,k} = \left(\overline{E}\right)^{\delta,k}$.

[25]This follows directly from the following two facts: (i) the Mertens-Zamir isomorphism $\mu_i : \mathcal{T}_i \to \Delta(\Theta \times \mathcal{T}_{-i})$ becomes a homeomorphism when each $\mathcal{T}_j$ is endowed with the product topology and $\Delta(\Theta \times \mathcal{T}_{-i})$ is endowed with

a formal sense in which convergence in common beliefs is the "common" analogue of product convergence.

Notice that common belief convergence implies product convergence, because setting $\mathbf{p} = (p, 0)$ yields $B_i^p = C_i^{\mathbf{p}}$. As for the connection with strategic convergence, we show:

Theorem 3. *Strategic convergence implies convergence in common beliefs.*

This result is useful for three reasons. First, it improves our understand of the strategic topology by providing an easy-to-interpret "lower bound" on how strong strategic convergence is. Second, the result (coupled with the example below), clarifies the relationship between the strategic topology for interim rationalizability on the universal type space and the ex ante strategic topology for BE in common prior, countable type spaces of the early literature on robustness. Third, the result is used in our proofs of the topological genericity results of Section 4).

As mentioned, the converse of Theorem 3 fails to hold:

Example 5 (Convergence in common beliefs does not imply strategic convergence)**.** We shall exhibit a sequence of types t_1^n that converges in common beliefs to a type t_1, but does not converge uniform-weakly on a frame; hence, by Theorem 2, it does not converge strategically either. To construct the sequence, fix $\theta_0 \in \Theta$ and $0 < p < q < 1$. For each player i, pick a type r_i that satisfies the following two conditions:

(i) r_i assigns probability zero to state θ_0;

(ii) for every product-closed proper subset $E \subset \Omega$ and $\eta > 0$, $r_i \notin C_i^{\eta \mathbf{1}}(E)$.[26]

Let s_i and t_i be the types who assign probability one to θ_0 and whose beliefs about the other player's types are as specified in Figure 1 below.

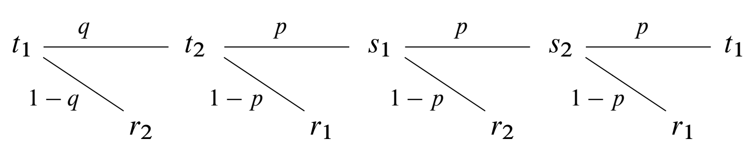

Figure 1: The types s_i and t_i.

the topology of weak convergence (as remarked in footnote 13); (ii) the Prohorov metric on $\Delta(\Theta \times \mathcal{T}_{-i})$ metrizes the topology of weak convergence.

[26]Types that satisfy these conditions exist. Ely and Pęski (2011, Theorem 1) show that the types that satisfy property (ii) are precisely those types to which product convergence is equivalent to strategic convergence, called *regular types*. They show that the set of regular types is a residual subset of the universal type space (in the product topology), in particular a non-empty set. As for condition (i), note that any type of player i that assigns probability one to some (θ, u_{-i}), where u_{-i} is a type of player $-i$ that satisfies property (ii), must also satisfy property (ii). This implies the existence of types satisfying both (i) and (ii).

Since r_1 and r_2 satisfy property (ii) above, for every $\mathbf{p} \in (0,1]^2$, every product-closed proper subset $E \subseteq \Omega$, and every $i \in I$,

$$t_i \in C_i^{\mathbf{p}}(E) \qquad \Longleftrightarrow \qquad \mathbf{p} \leqslant (p,p) \text{ and } E \supseteq \theta_0 \times \{s_1, t_1\} \times \{s_2, t_2\}. \tag{2}$$

In particular, no nontrivial event is common (q,p)-belief at t_1. We exploit this fact to construct the sequence t_1^n so that the probability assigned to (θ_0, t_2) drops to $q - \Delta \geqslant p$ under t_1^n, while ensuring that all events commonly $\mathbf{p}$-believed at t_1 for some $\mathbf{p} \in (0,1]^2$ remain so at t_1^n. Because the probability assigned to (θ_0, t_2) under t_1^n and t_1 differ by a positive Δ, we are able to construct a frame on which t_1^n fails to converge to t_1 uniform-weakly.

The construction of the sequence mimics the structure in Figure 1. Fix $0 < \Delta \leqslant q - p$ and for each player i define s_i^n and t_i^n as follows: let $t_1^1 = r_1$ and, for each $n \geqslant 1$, let s_2^n, s_1^n, t_2^n and t_1^{n+1} be the types who assign probability one to θ_0 and whose beliefs about the other player's types are as described in Figure 2 below.

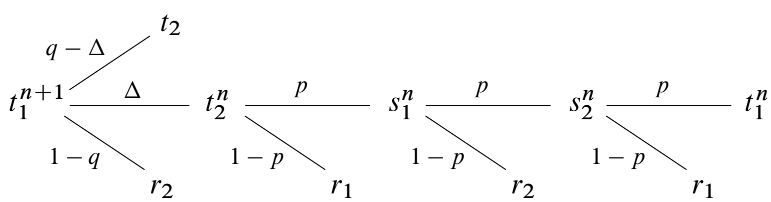

Figure 2: The sequences of types s_i^n and t_i^n.

The sequence t_1^n converges to t_1 in common beliefs. To see why, first note that $t_1^n \to t_1$ in the product topology: by the construction in Figures 1 and 2, s_2^1 has the same first-order belief as s_2, hence s_1^1 has the same second-order belief as s_1, which implies t_2^1 has the same third-order belief as t_2, and so forth. Second, given an arbitrary $\delta > 0$, by the construction in Figure 2 and the product-convergence $t_1^n \to t_1$, for each integer $k \geqslant 1$ we have $t_1^n \in B_1^{q-\Delta}\big(\theta_0 \times \{t_1\}^{\delta,k} \times t_2\big)$ for all n large enough. Since $t_2 \in C_2^{(p,p)}\big(\theta_0 \times \{s_1, t_1\} \times \{s_2, t_2\}\big)$ and $q - \Delta \geqslant p$, it follows that for all n large enough,

$$t_1^n \in B_1^p\Big(\theta_0 \times \{t_1\}^{\delta,k} \times C_2^{(p,p)}\big(\theta_0 \times \{s_1, t_1\} \times \{s_2, t_2\}\big)\Big) \subseteq C_1^{(p,p)}\Big(\big(\theta_0 \times \{s_1, t_1\} \times \{s_2, t_2\}\big)^{\delta,k}\Big).$$

It follows, by (2) and footnote 24, that $t_1^n \to t_1$ in common beliefs, as claimed.

To conclude the example, it remains to show:

Claim 1. There is a frame $\mathscr{P}$ such that $t_1^n \not\to t_1$ uniform-weakly on $\mathscr{P}$.

This claim is proved in Appendix A.5. Here, to provide intuition, we give a proof of the weaker but closely related statement that t_1^n does not converge to t_1 uniform-weakly. In effect, since t_1^1 assigns probability zero to θ_0, we have $d_1^1(t_1^1, t_1) = 1$. Then, by the constructions in Figures 1 and

2, the fourth-order Prohorov distance between t_2^1 and t_2 must also be one, that is, $d_2^4(t_2^1, t_2) = 1$. Therefore,

$$\mu_1\big((\theta_0 \times t_2)^{\Delta,4}\big|\, t_1^2\big) = \mu_1\big(\theta_0 \times t_2 \big|\, t_1^2\big) = q - \Delta = \mu_1(\theta_0 \times t_2|t_1) - \Delta,$$

hence $d_1^5(t_1^2, t_1) \geqslant \Delta$. Continuing in the same fashion it can be readily verified that $d_1^9(t_1^3, t_1) \geqslant \Delta$, $d_1^{13}(t_1^4, t_1) \geqslant \Delta$, and so forth. That is, $d_1^{4n-3}(t_1^n, t_1) \geqslant \Delta$ for all n, and hence $t_1^n \not\to t_1$ uniform-weakly, as claimed. ♦

This example, combined with Theorem 3, shows that strategic convergence is *strictly* stronger than convergence in common beliefs. However, when the limit type is assumed a *finite type*—a type that belongs to a finite type space—we show that convergence in common beliefs implies uniform-strategic convergence. We thus have:

Theorem 4. *Given a finite type t_i and a sequence of (possibly infinite) types t_i^n, the following statements are equivalent:*

(a) $t_i^n \to t_i$ *uniform-weakly;*

(b) $t_i^n \to t_i$ *uniform-strategically;*

(c) $t_i^n \to t_i$ *strategically;*

(d) $t_i^n \to t_i$ *uniform-weakly on every frame;*

(e) $t_i^n \to t_i$ *in common beliefs.*

4 Genericity Analysis

We apply our characterizations to establish topological genericity results on the universal type space concerning *critical types* and *common prior types*. We also provide analogous results for the measure-theoretic genericity notions of finite shyness and finite prevalence.[27]

4.1 Genericity of Critical Types

In a recent paper, Ely and Pęski (2011) define *critical types* as those types to which product convergence fails to imply strategic convergence. That is, a type t_i is critical if there is a sequence of types t_i^n that converges to t_i in the product topology such that, for some $\varepsilon > 0$ and some game

[27] Finite shyness strengthens the notion of shyness, originally proposed by Christensen (1974). See Anderson and Zame (2001) for details.

G, $R_i(t_i, G) \not\subseteq R_i(t_i^n, G, \varepsilon)$ infinitely often. Thus, from the point of view of robustness, the critical types are the problematic types.

How "large" is the set of critical types? Ely and Pęski (2011) provide two results that address this question. The first one is their insightful characterization of critical types in terms of common beliefs:

Ely and Pęski (2011, Theorem 1). *A type t_i is critical if and only if $t_i \in C_i^{p\mathbf{1}}(E)$ for some $p > 0$ and some product-closed proper subset $E_i \subset \mathcal{T}_i$.*

It follows that all finite types are critical, and so are almost all types that belong to a type space where there is a common prior (Ely and Pęski, 2011, Theorems 4 and 5). Thus, the set of critical types is "large" in the sense that critical types are pervasive: virtually all types used in applications are critical types.

An alternative approach to answer the question is based on topological genericity. There is, however, a tension between topological genericity in the product topology and the pervasiveness of critical types. Indeed, Ely and Pęski (2011) show that the set of critical types is topologically small (under the product topology):

Ely and Pęski (2011, Theorem 2). *The set of critical types is meager under the product topology on $\mathcal{T}_i$ (i.e., it is contained in a countable union of nowhere dense sets).*

We show, however, that this tension disappears when one considers either genericity under the strategic topology, or genericity under the measure-theoretic notion of finite prevalence:

Theorem 5. *The set of critical types is open and dense in the universal type space under the strategic topology. Furthermore, the set of critical types is finitely prevalent.*

Finally, consider the following variation of the notion of critical types: a type t_i is *uniformly critical* if there is some sequence that converges to t_i in the product topology but fails to converge uniform-strategically. An immediate implication of our Theorem 1 is that *all* types are uniformly critical, since for every type t_i in the universal type space there is always a sequence that converges to t_i in the product topology but does not converge uniform-weakly.[28]

[28]To prove this, for each player i and integer $k \geqslant 1$ we construct two measurable functions $\phi_i^k : \mathcal{T}_i \to \mathcal{T}_i$ and $\psi_i^k : \mathcal{T}_i \to \mathcal{T}_i$ that satisfy $d_i^{k-1}\big(t_i, \phi_i^k(t_i)\big) = d_i^{k-1}\big(t_i, \psi_i^k(t_i)\big) = 0$ for all $t_i \in \mathcal{T}_i$, and $\big(\phi_i^k(\mathcal{T}_i)\big)^{1,k} \cap \psi_i^k(\mathcal{T}_i) = \varnothing$. To define these functions for $k = 1$, pick any $\theta \in \Theta$ and $s_i, s_i' \in \mathcal{T}_i$ such that $\mu_i(\theta|s_i) = 1$ and $\mu_i(\theta|s_i') = 0$, and define $\phi_i^1(t_i) = s_i$ and $\psi_i^1(t_i) = s_i'$ for all $t_i \in \mathcal{T}_i$. Proceeding recursively, assume that the functions ϕ_i^k and ψ_i^k are defined for $k \geqslant 1$, and define ϕ^{k+1} and ψ_i^{k+1} as follows: for each $t_i \in \mathcal{T}_i$, $\theta \in \Theta$ and measurable $E_{-i} \subseteq \mathcal{T}_{-i}$, $\mu_i(\theta \times E_{-i}|\phi_i^{k+1}(t_i)) = \mu_i(\theta \times (\phi_{-i}^k)^{-1}(E_{-i})|t_i)$ and $\mu_i(\theta \times E_{-i}|\psi_i^{k+1}(t_i)) = \mu_i(\theta \times (\psi_{-i}^k)^{-1}(E_{-i})|t_i)$.

4.2 Genericity of Common Prior Types

The common prior assumption, according to which the players' beliefs are generated by a single probability distribution on the state space, is a cornerstone of virtually all models of information economics. Recall that a *common prior* on a countable type space $(T_i)_{i \in I}$ is a probability distribution $\eta \in \Delta(\Theta \times T_1 \times T_2)$ such that, for every $i \in I$, $t_i \in T_i$ and measurable $E \subseteq \Theta \times T_{-i}$,

$$\eta_i(t_i) = (\mathrm{marg}_{T_i}\, \eta)(t_i) > 0 \quad \text{and} \quad \eta\big(\big\{(\theta, t_1, t_2) : (\theta, t_{-i}) \in E\big\}\big)/\eta_i(t_i) = \mu_i(E|t_i).$$

A *common prior type* is a type that belongs to a countable type space that has a common prior.

The widespread use of common prior models, in both theoretical and applied work, begs the question of whether the behavioral implications of the common prior assumption are robust to misspecification errors in the assumed type space. Taking an interim perspective, Lipman (2003) shows that (finite) common prior types are dense in the product topology, but warns that this result should not be interpreted as a statement that the common prior assumption is without loss of generality: Although every type can be approximated by a common prior type in the product topology, the strategic behavior of that type can be very different from the strategic behavior of any approximating common prior type.

The next result shows that this lack of robustness is a pervasive phenomenon in the universal type space. The denseness result of Lipman (2003) is reversed, once we consider the strategic topology rather than the product topology. Moreover, an analogous conclusion holds when we look at measure-theoretic genericity.

Theorem 6. *The set of common prior types is nowhere dense in the universal type space under the strategic topology. Moreover, it is finitely shy.*

5 Conclusion

Going forward, we plan to extend our analysis to incorporate plausible restrictions on the class of games on which the strategic topology is based. Such restrictions may reflect the analyst's a priori knowledge of the payoff structure of the games he is interested in. He may, for example, know that the game of interest is a supermodular game or, say, a potential game. Another important restriction that is worth examining is motivated by mechanism design. On a fixed mechanism design environment (e.g., a single-unit auction environment with private values), the class of games that the designer can span by his choice of a mechanism does not generally span the whole space of games. We believe that an examination of the strategic impact of higher-order beliefs in restricted games is of paramount importance, as it brings our theoretical work closer to applications.

A Appendix

A.1 Properties of ICR

Some of the proofs in the appendix use the characterizations of ICR in terms of iterated elimination of never interim best-replies and of iterated elimination of strongly interim dominated strategies. We review these definitions below.

We begin with the recursive characterization of ICR in terms of iterated elimination of never interim best-replies. Given $\varepsilon \geqslant 0$ and a game $G = (A_i, g_i)_{i \in I}$, for each $i \in I$ and $t_i \in \mathcal{T}_i$, let $R_i^0(t_i, G, \varepsilon) = A_i$ and, recursively for each $k \geqslant 1$, let $R_i^k(t_i, G, \varepsilon)$ be the set of all $a_i \in A_i$ for which there is a measurable function $\sigma_{-i} : \Theta \times \mathcal{T}_{-i} \to \Delta(A_{-i})$ that satisfies:

$$\operatorname{supp} \sigma_{-i}(\theta, t_{-i}) \subseteq R_{-i}^{k-1}(t_{-i}, G, \varepsilon) \qquad \forall (\theta, t_{-i}) \in \Theta \times \mathcal{T}_{-i}, \tag{3}$$

$$\int_{\Theta \times \mathcal{T}_{-i}} \left[g_i(a_i, \sigma_{-i}(\theta, t_{-i}), \theta) - g_i(a_i', \sigma_{-i}(\theta, t_{-i}), \theta) \right] d\mu_i(\theta, t_{-i} | t_i) \geqslant -\varepsilon \qquad \forall a_i' \in A_i.$$

Then,

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

(See Dekel, Fudenberg, and Morris (2007), Corollary 1, Claim 2.)

For future reference, a measurable function σ_{-i} that satisfies (3) is called a $(k-1)$-*order ε-rationalizable conjecture*; a measurable function σ_{-i} such that $\operatorname{supp} \sigma_{-i}(\theta, t_{-i}) \subseteq R_{-i}(t_{-i}, G, \varepsilon)$ for every $(\theta, t_{-i}) \in \Theta \times \mathcal{T}_{-i}$ is called an *ε-rationalizable conjecture*.

Finally, we review the definition of ICR in terms of iterated elimination of strongly interim dominated actions. For each $i \in I$ and $t_i \in \mathcal{T}_i$, let $S_i^0(t_i, G, \varepsilon) = A_i$ and, recursively for $k \geqslant 1$, define $S_i^k(t_i, G, \varepsilon)$ as the set of all $a_i \in A_i$ such that for every *mixed* deviation $\alpha_i \in \Delta(A_i)$ there is a conjecture $\sigma_{-i} : \Theta \times \mathcal{T}_{-i} \to \Delta(A_{-i})$ that satisfies:

$$\operatorname{supp} \sigma_{-i}(\theta, t_{-i}) \subseteq S_{-i}^{k-1}(t_{-i}, G, \varepsilon) \qquad \forall (\theta, t_{-i}) \in \Theta \times \mathcal{T}_{-i}, \quad \text{and}$$

$$\int_{\Theta \times \mathcal{T}_{-i}} \left[g_i(a_i, \sigma_{-i}(\theta, t_{-i}), \theta) - g_i(\alpha_i, \sigma_{-i}(\theta, t_{-i}), \theta) \right] d\mu_i(\theta, t_{-i} | t_i) \geqslant -\varepsilon. \tag{4}$$

Then,

$$R_i^k(t_i, G, \varepsilon) = S_i^k(t_i, G, \varepsilon),$$

and thus,

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

(See Chen, Di Tillio, Faingold, and Xiong (2010), Proposition 1.) Likewise, $a_i \in R_i(t_i, G, \varepsilon)$ if and only if, for every $\alpha_i \in \Delta(A_i)$, there is a rationalizable conjecture $\sigma_{-i} : \Theta \times \mathcal{T}_{-i} \to \Delta(A_i)$ that satisfies (4).

Remark 1. In the above formulations of ICR we have chosen to work with conjectures of the form $\sigma_{-i} : \Theta \times \mathcal{T}_{-i} \to \Delta(A_{-i})$, rather than with those of the form $\nu \in \Delta(\Theta \times \mathcal{T}_{-i} \times A_{-i})$, $\text{marg}_{\Theta \times \mathcal{T}_{-i}} \nu = \mu_i(t_i)$, as in the definition of ICR given in Section 2.1. These two kinds of conjetures are related by the disintegration formula:

$$\nu(\theta \times E_{-i} \times a_{-i}) = \int_{E_{-i}} \sigma_{-i}(\theta, t_{-i})[a_{-i}]\, \mu_i(\theta \times dt_{-i}|t_i),$$

for every $\theta \in \Theta$, $a_{-i} \in A_{-i}$ and measurable subset $E_{-i} \subseteq \mathcal{T}_{-i}$.[29]

A.2 Proof of Theorem 1

The if direction is proved in Chen, Di Tillio, Faingold, and Xiong (2010). The proof of the only if direction relies on Lemma 3, Corollary 1 and Lemma 4 below.

We begin with some useful definitions and notations. Given a game $G = (A_i, g_i)_{i \in I}$, say that an action $a_i^0 \in A_i$ is a *zero action* for player i if $g_i\left(a_i^0, a_{-i}, \theta\right) = 0$ for all θ and a_{-i}. An action profile $a^c = (a_i^c, a_{-i}^c)$ is a *coordination pair* if $g_i\left(a_i^c, a_{-i}^c, \theta\right) = \max g_i$ for all θ and i. (In particular, any action of player i that is part of a coordination pair is rationalizable for any type of player i.) Finally, we write $\pi_{t_i,\sigma_{-i}} \in \Delta\left(A_{-i} \times \Theta\right)$ to denote the belief of type t_i over the actions of player $-i$ and the payoff-relevant states, when he has a conjecture $\sigma_{-i} : \Theta \times \mathcal{T}_{-i} \to \Delta(A_{-i})$; i.e.,

$$\pi_{t_i,\sigma_{-i}}\left(a_{-i}, \theta\right) \overset{\text{def}}{=} \int_{\mathcal{T}_{-i}} \sigma_{-i}(\theta, t_{-i})\left[a_{-i}\right]\, \mu_i(\theta \times dt_{-i}|t_i) \quad \forall (a_{-i}, \theta) \in A_{-i} \times \Theta.$$

The only if part of the theorem is a direct implication of Lemma 4 below. Lemma 3 and Corollary 1 are intermediate results.

Lemma 3. *For every $\varepsilon > 0$, integer $k \geqslant 1$, player j and finite set of finite types $\{t_{j,1}, t_{j,2}, \ldots, t_{j,N}\} \subset \mathcal{T}_j$, there is a game $G = (A_i, g_i)_{i \in I}$ with payoffs in the interval $[-5, 3]$, and a set of actions $\{a_{j,1}^*, a_{j,2}^*, \ldots, a_{j,N}^*\} \subset A_j$, such that:*

(i) every player i has a zero action $a_i^0 \in A_i$;

(ii) there is a coordination pair $a^c \in A_1 \times A_2$ such that $a_j^c \notin \left\{a_{j,1}^*, a_{j,2}^*, ..., a_{j,N}^*\right\}$ and $g_i\left(a_i, a_{-i}^c, \theta\right) \geqslant -2$ for every $a_i \in A_i$, $\theta \in \Theta$ and $i \in I$.

(iii) $R_i\left(\cdot, G, \gamma\right) = R_i^k\left(\cdot, G, \gamma\right)$ for every $i \in I$ and $\gamma \in [0, \frac{1}{2})$;

[29] Given a ν that satisfies $\text{marg}_{\Theta \times \mathcal{T}_{-i}} \nu = \mu_i(t_i)$, the disintegration fomula only pins down σ_{-i} up to a set of $\mu_i(t_i)$-probability zero. But, outside this null set, we can set σ_{-i} equal to a measurable selection from the correspondence $R_{-i}^{k-1}(\cdot, G, \varepsilon)$, thus ensuring that (3) is satisfied everywhere provided $\nu(\Theta \times \text{graph } R_{-i}^{k-1}) = 1$ (as opposed to almost everywhere). The fact that such a measurable selection exists follows from the upper hemi-continuity of $R_{-i}^{k-1}(\cdot, G, \varepsilon)$ (in the product topology) and the Kuratowsky-Nyll-Nardzewski Selection Theorem.

(iv) for every $1 \leqslant n \leqslant N$*,* $a^*_{j,n} \in R_j\left(t_{j,n}, G\right)$*;*

(v) for every $1 \leqslant n \leqslant N$ *and* $s_j \in \mathcal{T}_j$ *with* $d^k_j\left(s_j, t_{j,n}\right) > \varepsilon$*,* $a^*_{j,n} \notin R_j\left(s_j, G, \frac{\varepsilon}{2}\right)$*;*

Proof. The proof is by induction on k. Consider first $k = 1$ and $j = 1$ ($j = 2$ can be similarly proved). Fix a finite set of player 1's types $\left\{t_{1,1}, t_{1,2}, ..., t_{1,N}\right\}$. Enumerate the nonempty subsets of Θ as E_1, E_2,..., E_L. For each $(n, \ell) \in \{1, 2, \ldots, N\} \times \{0, 1, \ldots, L\}$ consider the function $\phi_{n,\ell} : \Theta \to [-1, 1]$ described in the following table:

	$\theta \in E_\ell$	$\theta \notin E_\ell$
$\ell = 0$	0	0
$\ell \geqslant 1$	$-\left(1 - \mu_1\left(E_\ell \vert t_{1,n}\right)\right)$	$\mu_1\left(E_\ell \vert t_{1,n}\right)$

Thus, the functions $\phi_{n,\ell}$ define an auxiliary game between player 1 and Nature, where $\ell = 0$ is a safe bet for player 1, and $\ell \geqslant 1$ is a risky bet on the event $\theta \notin E_\ell$. The rewards of the risky bets are such that:

- any type that has the same first-order beliefs as type $t_{1,n}$ is exactly indifferent between $\ell = 0$ and any $\ell \geqslant 1$;
- any type whose first-order belief is different from that of $t_{1,n}$ strictly prefers some risky bet $\ell \geqslant 1$ than the safe bet $\ell = 0$.

We use the functions $\phi_{n,\ell}$ to construct a game $G = (A_i, g_i)_{i \in I}$ to prove our claim for $k = 1$. In this game,

$$A_1 = \left(\{1, 2, \ldots, N\} \times \{0, 1, \ldots, L\}\right) \dot{\cup} \left\{a^0_1, a^c_1\right\} \text{ and } A_2 = \{1, 2, \ldots, N\} \ \dot{\cup} \ \left\{a^0_2, a^c_2\right\}.$$

Player 1's payoffs are specified as follows:

- a^0_1 is a zero action for player 1;
- if player 1 chooses a^c_1, she gets 3 if player 2 chooses a^c_2, and gets 0 otherwise (regardless of θ);
- if player 1 chooses $(n, \ell) \in \{1, \ldots, N\} \times \{0, \ldots, L\}$ and the state is θ, she gets $\phi_{n,\ell}(\theta)$ if player 2 chooses n, and she gets $\phi_{n,\ell}(\theta) - 1$ if player 2 chooses any action different from n.

Player 2's payoffs are specified as follows:

- Player 2 gets 3 if (a^c_1, a^c_2) is chosen (regardless of θ), and gets 0 otherwise.

Thus, when the state is θ, we can draw the payoff matrix $g(\cdot,\cdot,\theta)$ as follows:

	1	2	$\cdots$	N	a_2^0	a_2^c
$(1,\ell)$	$\phi_{1,\ell}(\theta),0$	$\phi_{1,\ell}(\theta)-1,0$	$\cdots$	$\phi_{1,\ell}(\theta)-1,0$	$\phi_{1,\ell}(\theta)-1,0$	$\phi_{1,\ell}(\theta)-1,0$
$(2,\ell)$	$\phi_{2,\ell}(\theta)-1,0$	$\phi_{2,\ell}(\theta),0$	$\cdots$	$\phi_{2,\ell}(\theta)-1,0$	$\phi_{2,\ell}(\theta)-1,0$	$\phi_{2,\ell}(\theta)-1,0$
$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\vdots$	$\vdots$
(N,ℓ)	$\phi_{N,\ell}(\theta)-1,0$	$\phi_{N,\ell}(\theta)-1,0$	$\cdots$	$\phi_{N,\ell}(\theta),0$	$\phi_{N,\ell}(\theta)-1,0$	$\phi_{N,\ell}(\theta)-1,0$
a_1^0	$0,0$	$0,0$	$\cdots$	$0,0$	$0,0$	$0,0$
a_1^c	$0,0$	$0,0$	$\cdots$	$0,0$	$0,0$	$3,3$

Since $\phi_{n,\ell}(\theta)\in[-1,1]$ for all n, ℓ and θ, game G has payoffs bounded between -2 and 3. We claim that G, along with the actions $a_{1,n}^{*}\overset{\text{def}}{=}(n,0)$, $n=1,\dots,N$, satisfy properties (i)–(v) for $k=1$. Properties (i) and (ii) clearly hold. To prove (iii), first note that $R_2(\cdot,G,\gamma)=R_2^1(\cdot,G,\gamma)=A_2$ for all $\gamma\geqslant 0$. Indeed, the profile of correspondences $(\varsigma_i)_{i\in I}$, where $\varsigma_2(t_2)=A_2$ and $\varsigma_1(t_1)=\{a_1^0,a_1^c\}$ for all t_1 and t_2, has the γ-best reply property. It follows that $R_1(\cdot,G,\gamma)=R_1^1(\cdot,G,\gamma)$ for all $\gamma\geqslant 0$. Hence, (iii) holds.

It remains to prove (iv) and (v). First, $(n,0)$ is rationalizable for $t_{1,n}$, since given the conjecture that player 2 plays n, type $t_{1,n}$ gets 0 by playing (n,ℓ) for any ℓ:

$$\mu_1\left(E_\ell|t_{1,n}\right)\cdot\left(-\left(1-\mu_1\left(E_\ell|t_{1,n}\right)\right)\right)+\left(1-\mu_1\left(E_\ell|t_{1,n}\right)\right)\cdot\mu_1\left(E_\ell|t_{1,n}\right)=0,$$

and gets at most 0 by playing any action not in $\{(n,0),\dots,(n,L)\}$. Thus, (iv) holds for $a_{1,n}^{*}=(n,0)$. Second, consider any type s_1 with $d_1^1\left(s_1,t_{1,n}\right)>\varepsilon$. Then, there exists some $1\leqslant\ell\leqslant L$ such that $\mu_1\left(E_\ell|t_{1,n}\right)>\mu_1\left(E_\ell^{\varepsilon,0}|s_1\right)+\varepsilon=\mu_1\left(E_\ell|s_1\right)+\varepsilon$.[30] Then, given any conjecture about the behavior of player 2, the difference in expected payoffs between (n,ℓ) and $(n,0)$ for type s_1 is

$$\begin{aligned}\mu_1(E_\ell|s_1)\cdot\left(-\left(1-\mu_1(E_\ell|t_{1,n})\right)\right)+(1-\mu_1(E_\ell|s_1))\cdot\mu_1(E_\ell|t_{1,n})&\\ =\mu_1(E_\ell|t_{1,n})-\mu_1(E_\ell|s_1)>\varepsilon.&\end{aligned}$$

Hence, $a_{1,n}^{*}=(n,0)$ is not ε-rationalizable for type s_1, which proves (v).

We now prove our claim for $k+1$ assuming that it holds for k. Again, we assume $j=1$, and the proof for $j=2$ is similar. Let $t_{1,1},\dots,t_{1,N}$ be arbitrary finite types of player 1. Consider the finite set $T_2=\{t_{2,1},t_{2,2},\dots,t_{2,N'}\}$ of all types of player 2 that are assigned positive probability by some $t_{1,n}$, for $n=1,\dots,N$. By the induction hypothesis, we can find a game $G=(A_i,g_i)_{i\in I}$, a set of actions $\{a_{2,1}^{*},a_{2,2}^{*},\dots,a_{2,N'}^{*}\}\subset A_2$, and action profiles a^0 and a^c in G, that satisfy properties (i)–(v) relative to the finite set of finite types T_2.

Let $T_2^k=\{t_{2,1}^k,t_{2,2}^k,\dots,t_{2,N'}^k\}$ be the set of k-order beliefs of types in T_2. Enumerate the nonempty subsets of $\Theta\times T_2^k$ as E_1, E_2,..., E_L. For each $1\leqslant\ell\leqslant L$, define

$$F_\ell=\left\{\left(a_{2,n'}^{*},\theta\right):1\leqslant n'\leqslant N',\left(\theta,t_{2,n'}^k\right)\in E_\ell\right\}.$$

[30] The equality follows because $d^0\equiv 0$ and Θ is endowed with the discrete metric.

For each $(n,\ell) \in \{1,\ldots,N\} \times \{0,\ldots,L\}$, define a function $\phi_{n,\ell} : A_2 \times \Theta \to [-1,1]$ as in the following table:[31]

	$(a_2,\theta) \in F_\ell$	$(a_2,\theta) \notin F_\ell$
$\ell = 0$	0	0
$\ell \geqslant 1$	$-\left(1-\mu_1\left(E_\ell \vert t_{1,n}\right)\right)$	$\mu_1\left(E_\ell \vert t_{1,n}\right)$

We use the game G and the functions $\phi_{n,\ell}$ to define a new game $\bar{G} = \left(\bar{A}_i, \bar{g}_i\right)$ to prove our claim for $k+1$. In this game,

$$\bar{A}_1 = A_1 \dot{\cup} \left(\{1,\ldots,N\} \times \{0,1,\ldots,L\}\right) \dot{\cup} \left\{\bar{a}_1^0\right\} \text{ and } \bar{A}_2 = A_2 \times \{0,1,\ldots,N\}.$$

Player 1's payoffs are specified as follows (see also the following table):

- $\bar{a}_1^0$ is a zero action for player 1;
- if player 1 chooses $a_1 \in A_1$ and the state is θ, he gets $g_1(a_1,a_2,\theta)$ if player 2 chooses $(a_2,0) \in A_2 \times \{0\}$, and gets $g_1\left(a_1,a_2^c,\theta\right) - 3$ otherwise.
- if player 1 chooses (n,ℓ) and the state is θ, he gets $\frac{\phi_{n,\ell}\left(a_2^c,\theta\right)}{2} - 1$ if player 2 chooses $(a_2,0) \in A_2 \times \{0\}$; he gets $\phi_{n,\ell}(a_2,\theta)$ if player 2 chooses $(a_2,n) \in A_2 \times \{n\}$; and he gets $\phi_{n,\ell}\left(a_2^c,\theta\right) - 1$ if player 2 chooses $(a_2,m) \in A_2 \times \{m\}$ with $m \neq n$ and $m \neq 0$.

	$A_2 \times \{0\}$	$A_2 \times \{1\}$	$A_2 \times \{2\}$	$\cdots$	$A_2 \times \{N\}$
A_1	$g_1(a_1,a_2,\theta)$	$g_1\left(a_1,a_2^c,\theta\right) - 3$	$g_1\left(a_1,a_2^c,\theta\right) - 3$	$\cdots$	$g_1\left(a_1,a_2^c,\theta\right) - 3$
$(1,\ell)$	$\frac{\phi_{1,\ell}\left(a_2^c,\theta\right)}{2} - 1$	$\phi_{1,\ell}(a_2,\theta)$	$\phi_{1,\ell}\left(a_2^c,\theta\right) - 1$	$\cdots$	$\phi_{1,\ell}\left(a_2^c,\theta\right) - 1$
$(2,\ell)$	$\frac{\phi_{2,\ell}\left(a_2^c,\theta\right)}{2} - 1$	$\phi_{2,\ell}\left(a_2^c,\theta\right) - 1$	$\phi_{2,\ell}(a_2,\theta)$	$\cdots$	$\phi_{2,\ell}\left(a_2^c,\theta\right) - 1$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$
(N,ℓ)	$\frac{\phi_{N,\ell}\left(a_2^c,\theta\right)}{2} - 1$	$\phi_{N,\ell}\left(a_2^c,\theta\right) - 1$	$\phi_{N,\ell}\left(a_2^c,\theta\right) - 1$	$\cdots$	$\phi_{N,\ell}(a_2,\theta)$
$\bar{a}_1^0$	0	0	0	$\cdots$	0

Player 2's payoffs are specified as follows:

- If player 2 chooses (a_2,m), he gets $g_2(a_2,a_1,\theta)$ if player 1 chooses $a_1 \in A_1$, and he gets $g_2\left(a_2,a_1^c,\theta\right)$ otherwise.

[31] Recall that we identify any measurable subset $E \subseteq \mathcal{T}_2^k$ with the cylinder

$$\{t_2 \in \mathcal{T}_2 : \text{the } k\text{-order belief of } t_2 \text{ belongs to } E\}.$$

By the induction hypothesis, game G satisfies property (ii) and has payoffs in the interval $[-5, 3]$; it follows that game $\bar{G}$ also has payoffs in the interval $[-5, 3]$.

We now prove that game $\bar{G}$, along with the actions $a^*_{1,n} \stackrel{\text{def}}{=} (n, 0)$, $n = 1, \ldots, N$, satisfy properties (i)–(v) for $k+1$. First, (i) follows because $\bar{a}^0_1$ and (a^0_2, n) are zero actions for players 1 and 2, respectively. Second, (ii) is satisfied by the coordination pair $\left(a^c_1, \left(a^c_2, 0\right)\right)$: for any $\bar{a}_1 \in \bar{A}_1$ and $\theta \in \Theta$,

$$\bar{g}_1\left(\bar{a}_1, \left(a^c_2, 0\right), \theta\right) \geqslant \min\left\{g_1\left(a_1, a^c_2, \theta\right), \min_{n,\ell} \phi_{n,\ell}\left(a^c_2, \theta\right)/2 - 1, 0\right\} \geqslant -2;$$

moreover, $\bar{g}_2\left(\left(a_2, n\right), a^c_1, \theta\right) = g_2\left(a_2, a^c_1, \theta\right) \geqslant -2$ for any $(a_2, n) \in \bar{A}_2$ and $\theta \in \Theta$, and $\bar{g}_1\left(a^c_1, \left(a^c_2, 0\right), \theta\right) = \bar{g}_2\left(\left(a^c_2, 0\right), a^c_1, \theta\right) = 3$ for any $\theta \in \Theta$.

The proof of (iii)–(v) relies on the following claim, whose proof is postponed.

Claim 2. For every integer $r \geqslant 0$ and $\gamma \in [0, 1/2)$,

1. $R^r_1\left(\cdot, \bar{G}, \gamma\right) \cap A_1 = R^r_1\left(\cdot, G, \gamma\right)$;
2. $R^r_2\left(\cdot, \bar{G}, \gamma\right) = R^r_2\left(\cdot, G, \gamma\right) \times \{0, 1, 2, \ldots, N\}$.

We now prove that Claim 2 implies properties (iii)–(v).

<u>(iii)</u>: By the induction hypothesis, $R_2\left(\cdot, G, \gamma\right) = R^k_2\left(\cdot, G, \gamma\right)$. Thus, Claim 2 implies that $R_2\left(\cdot, \bar{G}, \gamma\right) = R^k_2\left(\cdot, \bar{G}, \gamma\right)$. This, in turn, implies $R_1\left(\cdot, \bar{G}, \gamma\right) = R^{k+1}_1\left(\cdot, \bar{G}, \gamma\right)$ and hence (iii).

<u>(iv)</u>: Given any $1 \leqslant n \leqslant N$, consider the conjecture $\sigma_2 : \Theta \times T_2 \to \Delta(\bar{A}_2)$ such that $\sigma_2\left(\theta, t_{2,n'}\right)\left[a^*_{2,n'}, n\right] = 1$ for each $n' = 1, \ldots N'$ and $\theta \in \Theta$. By part 1 of Claim 2 and the fact that G satisfies property (iv) (by the induction hypothesis), σ_2 is a rationalizable conjecture in $\bar{G}$. Moreover, given such a conjecture, $t_{1,n}$ gets an expected payoff of 0 by playing (n, ℓ) for any ℓ, and gets at most 0 by playing any action in $\bar{A}_1 \setminus \{(n, 0), \ldots, (n, L)\}$:

$$\begin{aligned}
&\pi_{t_{1,n},\sigma_2}\left(F_\ell\right) \cdot \left(-\left(1 - \mu_1\left(E_\ell | t_{1,n}\right)\right)\right) + \left(1 - \pi_{t_{1,n},\sigma_2}\left(F_\ell\right)\right) \cdot \mu_1\left(E_\ell | t_{1,n}\right) \\
&\qquad = \mu_1\left(E_\ell | t_{1,n}\right) \cdot \left(-\left(1 - \mu_1\left(E_\ell | t_{1,n}\right)\right)\right) + \left(1 - \mu_1\left(E_\ell | t_{1,n}\right)\right) \cdot \mu_1\left(E_\ell | t_{1,n}\right) = 0.
\end{aligned}$$

In particular, $a^*_{1,n} = (n, 0)$ is a best reply for $t_{1,n}$.

<u>(v)</u>: Fix $1 \leqslant n \leqslant N$ and consider any type s_1 with $d^{k+1}_1\left(s_1, t_{1,n}\right) > \varepsilon$. Then, there exists some $1 \leqslant \ell \leqslant L$ such that

$$\mu_1\left(E_\ell | t_{1,n}\right) > \mu_1\left(\left(E_\ell\right)^{\varepsilon,k} \middle| s_1\right) + \varepsilon.$$

It follows that, given any $\frac{\varepsilon}{2}$-rationalizable conjecture $\sigma_2 : \Theta \times \mathcal{T}_2 \to \Delta(\bar{A}_2)$, the difference in expected payoffs between actions (n, ℓ) and $(n, 0)$ for type s_1 is at least $\varepsilon/2$. To prove this, we consider two cases separately: when player 2 chooses $m = n$; and when player 2 chooses $m \neq n$.

First, conditional on player 2 choosing n, the expected payoff difference between actions (n,ℓ) and $(n,0)$ for type s_1, given an arbitrary $\frac{\varepsilon}{2}$-rationalizable conjecture σ_2, is

$$\begin{aligned}&\sum_{a_2,\theta}\phi_{n,\ell}(a_2,\theta)\pi_{s_1,\sigma_2}(a_2,\theta|n)\\&\qquad=\quad\pi_{s_1,\sigma_2}(F_\ell|n)\cdot\left(-\left(1-\mu_1\left(E_\ell|t_{1,n}\right)\right)\right)+\left(1-\pi_{s_1,\sigma_2}(F_\ell|n)\right)\cdot\mu_1\left(E_\ell|t_{1,n}\right).\end{aligned}$$

But, for any $(\theta,s_2)\in\Theta\times\mathcal{T}_2$ and $n'=1,\ldots,N'$,

$$\begin{aligned}\sigma_2(\theta,s_2)[a^*_{2,n'},n]>0\quad&\Rightarrow\quad(a^*_{2,n'},n)\in R_2\left(s_2,\bar{G},\frac{\varepsilon}{2}\right)&&\text{(since }\sigma_2\text{ is }\frac{\varepsilon}{2}\text{-rationalizable)}\\&\Rightarrow\quad a^*_{2,n'}\in R_2\left(s_2,G,\frac{\varepsilon}{2}\right)&&\text{(by Claim 2)}\\&\Rightarrow\quad d_2^k(s_2,t_{2,n'})\leqslant\varepsilon,&&\text{(by the induction hypothesis)}\end{aligned}$$

and thus, if $\pi_{s_1,\sigma_2}(n)>0$,

$$\pi_{s_1,\sigma_2}(F_\ell|n)\leqslant\mu_1\left((E_\ell)^{\varepsilon,k}|s_1\right),$$

which implies

$$\begin{aligned}&\sum_{a_2,\theta}\phi_{n,\ell}(a_2,\theta)\pi_{s_1,\sigma_2}(a_2,\theta|n)\\&\quad\geqslant\quad\mu_1\left((E_\ell)^{\varepsilon,k}|s_1\right)\cdot\left(-\left(1-\mu_1\left(E_\ell|t_{1,n}\right)\right)\right)+\left(1-\mu_1\left((E_\ell)^{\varepsilon,k}|s_1\right)\right)\cdot\mu_1\left(E_\ell|t_{1,n}\right)\\&\quad=\quad\mu_1\left(E_\ell|t_{1,n}\right)-\mu_1\left((E_\ell)^{\varepsilon,k}|s_1\right)>\varepsilon.\end{aligned}$$

Second, conditional on player 2 choosing $m\neq n$, the expected payoff difference between actions (n,ℓ) and $(n,0)$ for type s_1 (given any conjecture) is at least $\phi_{n,\ell}(a_2^c,\theta)/2=\mu_1(E_\ell|t_{1,n})/2>\varepsilon/2$. (This is because $a_2^c\neq a^*_{2,n'}$ for all n', and hence $(a_2^c,\theta)\notin F_\ell$ for every θ.) We have thus shown that, given any $\varepsilon/2$-rationalizable conjecture, and conditional on any choice of $m=0,\ldots,N$ by player 2 with $\pi_{s_1,\sigma_2}(m)>0$, type s_1 gains at least $\varepsilon/2$ by deviating from $(n,0)$ to (n,ℓ). Thus, he also gains $\varepsilon/2$ unconditionally on m, and hence property (v) follows.

To conclude the proof, it remains to prove Claim 2. We prove it by induction on $r\geqslant 0$. First, the claim is trivially true for $r=0$. We now consider $r\geqslant 1$, assume that the claim holds for any $0\leqslant r'<r$, and prove that it also holds for r.

$\underline{R_1^r(t_1,G,\gamma)\supset R_1^r\left(t_1,\bar{G},\gamma\right)\cap A_1}$: Let $a_1\in R_1^r\left(t_1,\bar{G},\gamma\right)\cap A_1$. Then, there is an $(r-1)$-order γ-rationalizable conjecture $\bar{\sigma}_2:\Theta\times\mathcal{T}_2\to\Delta\left(\bar{A}_2\right)$ in $\bar{G}$ such that for any $a_1'\in A_1$,

$$\int_{\Theta\times\mathcal{T}_2}d\mu_1(\theta,t_2|t_1)\sum_{(a_2,n)\in\bar{A}_2}\left[\bar{g}_1(a_1,(a_2,n),\theta)-\bar{g}_1(a_1',(a_2,n),\theta)\right]\bar{\sigma}_2(\theta,t_2)\left[(a_2,n)\right]\geqslant-\gamma.\tag{5}$$

Consider the mapping $\varphi_2 : \bar{A}_2 \to A_2$,

$$\varphi_2(a_2, n) = \begin{cases} a_2, & \text{if } n = 0; \\ a_2^c, & \text{if } n \neq 0. \end{cases}$$

Define σ_2 as the conjecture in G such that $\sigma_2(\theta, t_2)[a_2] \overset{\text{def}}{=} \bar{\sigma}_2(\theta, t_2)\left[\varphi_2^{-1}(a_2)\right]$ for any $(\theta, t_2) \in \Theta \times \mathcal{T}_2$ and $a_2 \in A_2$. Since $\bar{\sigma}_2$ is $(r-1)$-order γ-rationalizable in $\bar{G}$, $\bar{\sigma}_2(\theta, t_2)[(a_2, 0)] > 0$ implies $(a_2, 0) \in R_2^{r-1}\left(t_2, \bar{G}, \gamma\right)$, and by the induction hypothesis, $a_2 \in R_2^{r-1}(t_2, G, \gamma)$. Moreover, a_2^c is part of a coordination pair in G, hence a_2^c is rationalizable in G for any type. Thus, σ_2 is an $(r-1)$-order γ-rationalizable conjecture in G. Moreover, for any $a_1' \in A_1$,

$$\begin{aligned}
&\int_{\Theta\times\mathcal{T}_2} d\mu_1(\theta, t_2|t_1) \sum_{(a_2,n)\in\bar{A}_2} \left[\bar{g}_1(a_1, (a_2, n), \theta) - \bar{g}_1\left(a_1', (a_2, n), \theta\right)\right] \bar{\sigma}_2(\theta, t_2)[(a_2, n)] \\
= \quad &\int_{\Theta\times\mathcal{T}_2} d\mu_1(\theta, t_2|t_1)\Bigg(\sum_{a_2\in A_2} \left[g_1(a_1, a_2, \theta) - g_1\left(a_1', a_2, \theta\right)\right] \bar{\sigma}_2(\theta, t_2)[(a_2, 0)] \\
&\qquad + \left[\left(g_1\left(a_1, a_2^c, \theta\right) - 3\right) - \left(g_1\left(a_1', a_2^c, \theta\right) - 3\right)\right] \bar{\sigma}_2(\theta, t_2)\left[\left\{(a_2, n) \in \bar{A}_2 : n > 0\right\}\right]\Bigg) \\
= &\int_{\Theta\times\mathcal{T}_2} d\mu_1(\theta, t_2|t_1) \sum_{a_2\in A_2} \left[g_1(a_1, a_2, \theta) - g_1\left(a_1', a_2, \theta\right)\right] \sigma_2(\theta, t_2)[a_2] \qquad (6)
\end{aligned}$$

Then, (5) and (6) imply $a_1 \in R_1^r(t_1, G, \gamma)$.

$\underline{R_1^r(t_1, G, \gamma) \subset R_1^r\left(t_1, \bar{G}, \gamma\right) \cap A_1}$: Let $a_1 \in R_1^r(t_1, G, \gamma)$. Then, there exists an $(r-1)$-order γ-rationalizable conjecture $\sigma_2 : \Theta \times \mathcal{T}_2 \to \Delta(A_2)$ in G such that for any $a_1' \in A_1$,

$$\int_{\Theta\times\mathcal{T}_2} d\mu_1(\theta, t_2|t_1) \sum_{a_2\in A_2} \left[g_1(a_1, a_2, \theta) - g_1\left(a_1', a_2, \theta\right)\right] \sigma_2(\theta, t_2)[a_2] \geqslant -\gamma. \qquad (7)$$

Define $\bar{\sigma}_2$ as the conjecture in $\bar{G}$ such that $\bar{\sigma}_2(\theta, t_2)[(a_2, 0)] \overset{\text{def}}{=} \sigma_2(\theta, t_2)[a_2]$ for any $(\theta, t_2) \in \Theta \times \mathcal{T}_2$ and $a_2 \in A_2$ (and thus $\bar{\sigma}_2(\theta, t_2)[(a_2, n)] = 0$ for any $n > 0$). Since σ_2 is $(r-1)$-order γ-rationalizable in G , $\sigma_2(\theta, t_2)[a_2] > 0$ implies $a_2 \in R_2^{r-1}(t_2, G, \gamma)$, and by the induction hypothesis, $(a_2, 0) \in R_2^{r-1}\left(t_2, \bar{G}, \gamma\right)$. Hence, $\bar{\sigma}_2$ is $(r-1)$-order γ-rationalizable in G. We will now show that a_1 is a γ-best reply to $\bar{\sigma}_2$ for t_1 in $\bar{G}$. First, by (7) and the definition of $\bar{\sigma}_2$,

$$\int_{\Theta\times\mathcal{T}_2} d\mu_1(\theta, t_2|t_1) \sum_{a_2\in A_2} \left[g_1(a_1, a_2, \theta) - g_1\left(a_1', a_2, \theta\right)\right] \bar{\sigma}_2(\theta, t_2)[(a_2, 0)] \geqslant -\gamma \quad \forall a_1' \in A_1. \qquad (8)$$

Second, setting $a_1' = a_1^0$ in (7) and recalling that $\gamma < 1/2$,

$$\int_{\Theta\times\mathcal{T}_2} d\mu_1(\theta, t_2|t_1) \sum_{a_2\in A_2} g_1(a_1, a_2, \theta) \sigma_2(\theta, t_2)[a_2] \geqslant -\gamma > -1/2.$$

Since $\phi_{n,\ell}\left(a_2^c,\theta\right)/2-1\leqslant -1/2$ for any n and ℓ,

$$\int_{\Theta\times\mathcal{T}_2} d\mu_1(\theta,t_2|t_1)\sum_{a_2\in A_2}\left[\bar{g}_1\left(a_1,a_2,\theta\right)-\bar{g}_1\left((n,\ell),a_2,\theta\right)\right]\bar{\sigma}_2\left(\theta,t_2\right)\left[(a_2,0)\right]\geqslant 0 \quad \forall n,\ell. \tag{9}$$

By (8) and (9), a_1 is a γ-best reply to $\bar{\sigma}_2$ for t_1 in $\bar{G}$, and hence $a_1\in R_1^r\left(t_1,\bar{G},\gamma\right)$.

$R_2^r\left(t_2,\bar{G},\gamma\right)\subset R_2^r\left(t_2,G,\gamma\right)\times\{0,1,\ldots,N\}$: Suppose $(a_2,m)\in R_2^r\left(t_2,\bar{G},\gamma\right)$. Then, there is an $(r-1)$-order γ-rationalizable conjecture $\bar{\sigma}_1:\Theta\times\mathcal{T}_1\to\Delta\left(\bar{A}_1\right)$ in $\bar{G}$ such that for any $\left(a_2',m'\right)\in\bar{A}_2$,

$$\int_{\Theta\times\mathcal{T}_1} d\mu_2(\theta,t_1|t_2)\sum_{\bar{a}_1\in\bar{A}_1}\left[\bar{g}_2\left(\left(a_2,m\right),\bar{a}_1,\theta\right)-\bar{g}_2\left(\left(a_2',m'\right),\bar{a}_1,\theta\right)\right]\bar{\sigma}_1\left(\theta,t_1\right)\left[\bar{a}_1\right]\geqslant -\gamma.$$

Consider the map $\varphi_1:\bar{A}_1\to A_1$,

$$\varphi_1\left(\bar{a}_1\right)=\begin{cases}\bar{a}_1, & \text{if } \bar{a}_1\in A_1;\\ a_1^c, & \text{if } \bar{a}_1\notin A_1.\end{cases}$$

Let σ_1 be the conjecture in G such that $\sigma_1\left(\theta,t_1\right)\left[a_1\right]\overset{\text{def}}{=}\bar{\sigma}_1\left(\theta,t_1\right)\left[\varphi_1^{-1}\left(a_1\right)\right]$ for any $(\theta,t_1)\in\Theta\times\mathcal{T}_1$ and $a_1\in A_1$. Since $\bar{\sigma}_1$ is $(r-1)$-order γ-rationalizable in $\bar{G}$, for any $\bar{a}_1\in A_1$, $\bar{\sigma}_1\left(\theta,t_1\right)\left[\bar{a}_1\right]>0$ implies $\bar{a}_1\in R_1^{r-1}\left(t_1,\bar{G},\gamma\right)$, and by the induction hypothesis, $\bar{a}_1\in R_1^{r-1}\left(t_1,G,\gamma\right)$. Moreover, a_1^c is part of a coordination pair in G, hence it is rationalizable for any type. Thus, σ_1 is an $(r-1)$-order γ-rationalizable conjecture in G. Moreover, (5) implies that for any $a_2'\in A_2$,

$$\begin{aligned}
-\gamma &\leqslant \int_{\Theta\times\mathcal{T}_1} d\mu_2(\theta,t_1|t_2)\sum_{\bar{a}_1\in\bar{A}_1}\left[\bar{g}_2\left(\left(a_2,m\right),\bar{a}_1,\theta\right)-\bar{g}_2\left(\left(a_2',m'\right),\bar{a}_1,\theta\right)\right]\bar{\sigma}_1\left(\theta,t_1\right)\left[\bar{a}_1\right]\\
&= \int_{\Theta\times\mathcal{T}_1} d\mu_2(\theta,t_1|t_2)\Bigg(\sum_{a_1\in A_1}\left[g_2\left(a_2,a_1,\theta\right)-g_2\left(a_2',a_1,\theta\right)\right]\bar{\sigma}_1\left(\theta,t_1\right)\left[a_1\right]\\
&\quad+\left[g_2\left(a_2,a_1^c,\theta\right)-g_2\left(a_2',a_1^c,\theta\right)\right]\bar{\sigma}_1\left(\theta,t_1\right)\left[\bar{A}_1\backslash A_1\right]\Bigg)\\
&= \int_{\Theta\times T_1} d\mu_2(\theta,t_1|t_2)\sum_{a_1\in A_1}\left[g_2\left(a_2,a_1,\theta\right)-g_2\left(a_2',a_1,\theta\right)\right]\sigma_1\left(\theta,t_1\right)\left[a_1\right]
\end{aligned}$$

Therefore, $a_2\in R_2^r\left(t_2,G,\gamma\right)$.

$R_2^r\left(t_2,\bar{G},\gamma\right)\supset R_2^r\left(t_2,G,\gamma\right)\times\{0,1,\ldots,N\}$: Let $(a_2,m)\in R_2^r\left(t_2,G,\gamma\right)\times\{0,1,\ldots,N\}$. Then, there is an $(r-1)$-order γ-rationalizable conjecture $\sigma_1:\Theta\times\mathcal{T}_1\to\Delta\left(A_1\right)$ in G such that for any $a_2'\in A_2$,

$$\int_{\Theta\times\mathcal{T}_1} d\mu_2(\theta,t_1|t_2)\sum_{a_1\in A_1}\left[g_2\left(a_2,a_1,\theta\right)-g_2\left(a_2',a_1,\theta\right)\right]\sigma_1\left(\theta,t_1\right)\left[a_1\right]\geqslant -\gamma. \tag{10}$$

Let $\bar{\sigma}_1$ be the conjecture in $\bar{G}$ such that $\bar{\sigma}_1(\theta, t_1)[a_1] \stackrel{\text{def}}{=} \sigma_1(\theta, t_1)[a_1]$ for any $(\theta, t_1) \in \Theta \times \mathcal{T}_1$ and $a_1 \in A_1$ (and thus $\bar{\sigma}_1(\theta, t_1)[a_1] = 0$ for any $a_1 \notin A_1$). Since σ_1 is $(r-1)$-order γ-rationalizable in G, $\sigma_1(\theta, t_1)[a_1] > 0$ implies $a_1 \in R_1^{r-1}(t_1, G, \gamma)$, and by the induction hypothesis, $a_1 \in R_1^{r-1}\left(t_1, \bar{G}, \gamma\right)$. Hence, $\bar{\sigma}_1$ is $(r-1)$-order γ-rationalizable in $\bar{G}$. Since $\bar{g}_2((a_2, m), a_1, \theta) = \bar{g}_2((a_2, m'), a_1, \theta)$ for any m and m', it follows from (10) that $(a_2, m) \in R_2^r\left(t_2, \bar{G}, \gamma\right)$. ■

Corollary 1. *For every $\varepsilon > 0$, player i, positive integer k and finite type $t_i \in \mathcal{T}_i$, there exists a game G with $|G| \leqslant M$, and an action a_i of player i in G, such that:*

(i) $R_j(\cdot, G, \gamma) = R_j^k(\cdot, G, \gamma)$ *for every* $\gamma \in [0, M/10)$ *and* $j \in I$*;*

(ii) $a_i \in R_i(t_i, G)$*;*

(iii) $a_i \notin R_i^k(s_i, G, M\varepsilon/10)$ *for every* $s_i \in \mathcal{T}_i$ *with* $d_i^k(s_i, t_i) > \varepsilon$.

Proof. Immediate implication of Lemma 3, upon rescaling the payoffs by a factor of $M/5$. ■

Lemma 4. *For every $\varepsilon > 0$ there exists $\delta > 0$ such that for every $i \in I$ and s_i, $t_i \in \mathcal{T}_i$ with $d_i^{UW}(s_i, t_i) > \varepsilon$ there is a game G with $|G| \leqslant M$ such that $R_i(t_i, G) \not\subseteq R_i(s_i, G, \delta)$.*

Proof. Fix an $\varepsilon > 0$, a player i, an integer $k \geqslant 1$ and types s_i, $t_i \in \mathcal{T}_i$ with $d_i^k(s_i, t_i) > \varepsilon$. Fix $0 < \delta < M\varepsilon/10$ and choose $\rho > 0$ small enough that

$$\frac{M(\varepsilon - \rho)}{10} - 4M\rho > \delta. \tag{11}$$

Since finite types are dense in the product topology, there is a finite type t_i' such that $d_i^k\left(t_i', t_i\right) < \rho$. Then, $d_i^k\left(s_i, t_i'\right) > \varepsilon - \rho$. By Corollary 1, there is some game $G' = (A_i', g_i')_{i \in I}$ with $|G'| \leqslant M$, and some action a_i' of player i in G', such that $a_i' \in R_i\left(t_i', G', 0\right)$ and $a_i' \notin R_i^k\left(s_i, G', \frac{M(\varepsilon - \rho)}{10}\right)$. By Proposition 2 in Chen, Di Tillio, Faingold, and Xiong (2010), $a_i' \in R_i^k(t_i, G', 4M\rho)$. Then, it follows from (i) of Corollary 1 that $a_i' \in R_i(t_i, G', 4M\rho)$ and $a_i' \notin R_i\left(s_i, G', \frac{M(\varepsilon - \rho)}{10}\right)$.

To conclude, consider the game $G = (A_j, g_j)_{j \in I}$, defined as follows:

$$A_i = A_i', \quad A_{-i} = A_{-i}' \times A_i',$$

$$g_i(a_i, (a_{-i}^{-i}, a_{-i}^i), \theta) = \begin{cases} g_i'(a_i, a_{-i}^{-i}, \theta) + 4M\rho & : \quad a_i^i = a_{-i}^i \\ g_i'(a_i, a_{-i}^{-i}, \theta) & : \quad \text{otherwise,} \end{cases}$$

and

$$g_{-i}((a_{-i}^{-i}, a_{-i}^i), a_i, \theta) = g_{-i}'(a_{-i}^{-i}, a_i, \theta).$$

In game G, player $-i$ is indifferent among all actions a_{-i}^i; moreover, player i gets an additional payoff of $4M\rho$ whenever his action matches player $-i$'s choice of a_{-i}^i. Therefore, $R_i(\cdot, G, \gamma) = R_i(\cdot, G', \gamma + 4M\rho)$ for every $\gamma \geqslant 0$. In particular, we have that $a_i' \in R_i(t_i, G, 0)$ and $a_i' \notin R_i\left(s_i, G, \frac{M(\varepsilon - \rho)}{10} - 4M\rho\right) \supseteq R_i(s_i, G, \delta)$, where the inclusion follows from (11).[32] ■

[32] If necessary, rescale the payoffs from G to ensure $|G| \leqslant M$, and rescale δ by the same factor.

The only if direction of Theorem 1 then follows directly from Lemma 4.

A.3 Proof of Theorem 2

We begin with the following auxiliary result about the structure of ICR.

Lemma 5. *Fix a game $G = (A_i, g_i)_{i \in I}$ and $\varepsilon \geqslant 0$. For every integer $k \geqslant 1$, $i \in I$ and $t_i \in \mathcal{T}_i$, we have $a_i \in R_i^k(t_i, G, \varepsilon)$ if and only if, for every $\alpha_i \in \Delta(A_i)$,*

$$\sum_{\theta \in \Theta, B \subseteq A_{-i}} \max_{a_{-i} \in B} \Big[g_i(a_i, a_{-i}, \theta) - g_i(\alpha_i, a_{-i}, \theta) \Big] \times \mu_i\Big(\theta \times \Big\{ t_{-i} : R_{-i}^{k-1}(t_{-i}, G, \varepsilon) = B \Big\} \Big| t_i \Big) \geqslant -\varepsilon. \quad (12)$$

Likewise, $a_i \in R_i(t_i, G, \varepsilon)$ if and only if, for every $\alpha_i \in \Delta(A_i)$,

$$\sum_{\theta \in \Theta, B \subseteq A_{-i}} \max_{a_{-i} \in B} \Big[g_i(a_i, a_{-i}, \theta) - g_i(\alpha_i, a_{-i}, \theta) \Big] \times \mu_i\Big(\theta \times \Big\{ t_{-i} : R_{-i}(t_{-i}, G, \varepsilon) = B \Big\} \Big| t_i \Big) \geqslant -\varepsilon. \quad (13)$$

Proof. Straightforward implication of the characterization of ICR in terms of iterated dominance (see Appendix A.1). ■

The if direction of the theorem is an immediate consequence of the following lemma:

Lemma 6. *Fix $\delta > 0$ and a game $G = (A_i, g_i)_{i \in I}$ and let $\mathcal{P}$ denote the strategic frame associated with G. For every integer $k \geqslant 0$, $i \in I$ and $s_i, t_i \in \mathcal{T}_i$,*

$$d_{i,\mathcal{P}}^k(s_i, t_i) < \delta \quad \Longrightarrow \quad R_i(t_i, G) \subseteq R_i^k(s_i, G, 4M\delta).$$

In particular, for every $i \in I$ and s_i, $t_i \in \mathcal{T}_i$,

$$d_{i,\mathcal{P}}^{UW}(s_i, t_i) < \delta \quad \Longrightarrow \quad R_i(t_i, G) \subseteq R_i(s_i, G, 4M\delta).$$

Proof. We need only prove the first result, as the second result is a straightforward implication of the first one. For $k = 0$ the result is trivially true, as $R_i^0 \equiv A_i$. Proceeding by induction, we assume the result is true for $k \geqslant 0$ and show that it remains true for $k + 1$. Consider a pair of types s_i, t_i with $d_{i,\mathcal{P}}^{k+1}(s_i, t_i) < \delta$. Fix an arbitrary $a_i \in R_i(t_i, G)$ and let us show that $a_i \in R_i^{k+1}(s_i, G, 4M\delta)$. Given $\alpha_i \in \Delta(A_i)$, by Lemma 5 we have

$$\sum_{\theta \in \Theta, B \subseteq A_{-i}} \Delta g_i\big(\theta, B\big)\, \mu_i\big(\theta \times [B] \big| t_i\big) \geqslant 0, \quad (14)$$

where, for each $\theta \in \Theta$ and nonempty $B \subseteq A_{-i}$,

$$[B] \overset{\text{def}}{=} \Big\{ t_{-i} : R_{-i}(t_{-i}, G) = B \Big\},$$
$$\Delta g_i(\theta, B) \overset{\text{def}}{=} \max_{a_{-i} \in B} \Big[g_i(a_i, a_{-i}, \theta) - g_i(\alpha_i, a_{-i}, \theta) \Big].$$

By Lemma 5, in order to prove that $a_i \in R_i^{k+1}(s_i, G, 4M\delta)$ we need only show that

$$\sum_{\theta \in \Theta, B \subseteq A_{-i}} \Delta g_i\big(\theta, B\big)\, \mu_i\big(\theta \times [[B]] \big| s_i\big) \geqslant -4M\delta,$$

where

$$[[B]] \overset{\text{def}}{=} \Big\{ t_{-i} : R_{-i}^k\big(t_{-i}, G, 4M\delta\big) = B \Big\} \qquad \forall B \subseteq A_{-i}.$$

To prove this, first note that the induction hypothesis implies

$$[B]_{\mathcal{P}}^{\delta,k} \subseteq \bigcup_{C \supseteq B} [[C]] \qquad \forall B \subseteq A_{-i}. \tag{15}$$

Second, enumerate the elements of the finite set $\Theta \times \big\{ B : \varnothing \neq B \subseteq A_{-i} \big\}$ as $\big\{(\theta_n, B_n)\big\}_{n=1}^{N}$ $(N = \#\Theta(2^{\# A_{-i}} - 1))$ so that

$$\Delta g_i(\theta_n, B_n) \geqslant \Delta g_i(\theta_{n+1}, B_{n+1}) \qquad \forall n = 1, \ldots, N-1 \tag{16}$$

and

$$(\theta_m = \theta_n \quad \text{and} \quad B_m \supseteq B_n) \quad \Longrightarrow \quad m \leqslant n.\text{[33]} \tag{17}$$

[33] Such enumeration is possible because $\Delta g_i(\theta, B') \geqslant \Delta g_i(\theta, B)$ whenever $B' \supseteq B$.

Then,

$$
\begin{aligned}
&\sum_{\theta\in\Theta, B\subseteq A_{-i}} \Delta g_i\left(\theta, B\right)\mu_i\left(\theta\times[[B]]\middle|s_i\right) \\
&\quad\geqslant \sum_{n=1}^{N} \Delta g_i\left(\theta_n, B_n\right)\Big(\mu_i\left(\theta_n\times[[B_n]]\middle|s_i\right)-\mu_i\left(\theta_n\times[B_n]\middle|t_i\right)\Big) && (18)\\
&\quad= \sum_{n=1}^{N-1}\Big(\Delta g_i\left(\theta_n, B_n\right)-\Delta g_i\left(\theta_{n+1}, B_{n+1}\right)\Big) \\
&\qquad\qquad\times\sum_{m=1}^{n}\left(\mu_i\left(\theta_m\times[[B_m]]\middle|s_i\right)-\mu_i\left(\theta_m\times[B_m]\middle|t_i\right)\right) && (19)\\
&\quad= \sum_{n=1}^{N-1}\Big(\Delta g_i\left(\theta_n, B_n\right)-\Delta g_i\left(\theta_{n+1}, B_{n+1}\right)\Big) \\
&\qquad\qquad\times\Bigg(\mu_i\Big(\bigcup_{m=1}^{n}\theta_m\times[[B_m]]\Big|s_i\Big)-\mu_i\Big(\bigcup_{m=1}^{n}\theta_m\times[B_m]\Big|t_i\Big)\Bigg) && (20)\\
&\quad\geqslant \sum_{n=1}^{N-1}\Big(\Delta g_i\left(\theta_n, B_n\right)-\Delta g_i\left(\theta_{n+1}, B_{n+1}\right)\Big) \\
&\qquad\qquad\times\Bigg(\mu_i\Big(\Big(\bigcup_{m=1}^{n}\theta_m\times[B_m]\Big)^{\delta,k}_{\mathcal{P}}\Big|s_i\Big)-\mu_i\Big(\bigcup_{m=1}^{n}\theta_m\times[B_m]\Big|t_i\Big)\Bigg) && (21)\\
&\quad\geqslant \sum_{n=1}^{N-1}\left(\Delta g_i\left(\theta_n, B_n\right)-\Delta g_i\left(\theta_{n+1}, B_{n+1}\right)\right)\left(-\delta\right) && (22)\\
&\quad= -\left(\Delta g_i\left(\theta_1, B_1\right)-\Delta g_i\left(\theta_N, B_N\right)\right)\delta \\
&\quad\geqslant -4M\delta, && (23)
\end{aligned}
$$

where (18) follows from (14); (19) follows by a standard "integration by parts" argument; (20) follows from the fact that $\left\{(\theta_m, [B_m])\right\}_{m=1}^{N}$ and $\left\{(\theta_m, [[B_m]])\right\}_{m=1}^{N}$ are partitions of $\Theta\times\mathcal{T}_{-i}$; (21) follows from (15), (16) and (17); (22) follows from the assumption that $d^{k+1}_{i,\mathcal{P}}(s_i, t_i)\leqslant\delta$, the fact that $[B_m]\in\mathcal{P}_{-i}$ and (16); and (23) follows from $|\Delta g_i|\leqslant 2M$. ■

Before turning to the proof of the only if direction of Theorem 2, we introduce some notation. First, given a frame $\mathcal{P}$ and a type t_i, let $\mu_i(t_i)\big|_{\Theta\times\mathcal{P}_{-i}}$ denote the belief of type t_i over the events in $\Theta\times\mathcal{P}_{-i}$. We thus view $\mu_i(t_i)\big|_{\Theta\times\mathcal{P}_{-i}}$ as an element of the finite-dimensional simplex $\Delta(\Theta\times\mathcal{P}_{-i})$ (viewed as a subset of the Euclidean space $\mathbb{R}^{\#\Theta\cdot\#\mathcal{P}_{-i}}$). Second, let $\varphi_i : \Delta(\Theta\times\mathcal{P}_{-i})\to\mathcal{P}_i$ designate the function that maps each $q\in\Delta(\Theta\times\mathcal{P}_{-i})$ into $\mathcal{P}_i(t_i)$, where t_i is some type with $\mu_i(t_i)\big|_{\Theta\times\mathcal{P}_{-i}} = q$. Since $\mathcal{P}$ is a frame, the definition of φ_i is independent of the choice of t_i. Third, given a game $G=(A_i, g_i)_{i\in I}$, for each $\pi_i\in\Delta(A_i\times\Theta)$ define a probability distribution

$\pi_i^{\mathcal{P}} \in \Delta(\Theta \times \mathcal{P}_i)$ as follows:

$$\pi_i^{\mathcal{P}}(\theta, E) \stackrel{\text{def}}{=} \sum_{a_i \in \varphi_i^{-1}(E)} \pi_i(a_i, \theta) \qquad \forall (\theta, E) \in \Theta \times \mathcal{P}_i.$$

The proof of the only if direction of Theorem 2 relies on Lemmas 7 and 8 and Corollary 2 below.

Lemma 7. *Fix a frame $\mathcal{P}$. For each $\zeta > 0$ there exist $\varepsilon > 0$ and a game $G = (A_i, g_i)_{i \in I}$ such that for each $i \in I$,*

(i) A_i is a finite subset of $\Delta(\Theta \times \mathcal{P}_{-i})$ that is $\sqrt{\varepsilon}$-dense (relative to the Euclidean norm $\|\cdot\|$) in every element of the partition $\{\varphi_i^{-1}(E) : E \in \mathcal{P}_i\}$;

(ii) for each $a_{-i}, a'_{-i} \in A_{-i}$,

$$\varphi_{-i}(a_{-i}) = \varphi_{-i}(a'_{-i}) \quad \Longrightarrow \quad g_i(a_i, a_{-i}, \theta) = g_i(a_i, a'_{-i}, \theta) \quad \forall \theta \in \Theta, a_i \in A_i;$$

(iii) for each $\pi_{-i} \in \Delta(A_{-i} \times \Theta)$,

$$BR_i(\pi_{-i}, G, \varepsilon) \supseteq \left\{a_i \in A_i : \|a_i - \pi_{-i}^{\mathcal{P}}\| < \sqrt{\varepsilon}\right\};$$

(iv) for each $\pi_{-i}, \pi'_{-i} \in \Delta(A_{-i} \times \Theta)$ with $\|\pi_{-i}^{\mathcal{P}} - \pi'^{\mathcal{P}}_{-i}\| > \zeta$,

$$BR_i(\pi'_{-i}, G, 2\varepsilon) \cap \left\{a_i \in A_i : \|a_i - \pi_{-i}^{\mathcal{P}}\| < \sqrt{\varepsilon}\right\} = \varnothing.$$

Proof. Fix $\zeta > 0$ and a frame $\mathcal{P}$. Let $0 < \varepsilon < \zeta^2/(1 + \sqrt{3})^2$. Cover the finite-dimensional simplex $\Delta(\Theta \times \mathcal{P}_{-i})$ by a finite union of open balls $B_1, \ldots, B_N$ of diameter $\sqrt{\varepsilon}$. Select one point from $B_n \cap \varphi_i^{-1}(E)$, for each $n = 1, \ldots, N$ and $E \in \mathcal{P}_i$, and let A_i denote the set of selected points. By construction, A_i satisfies (i).

Consider the quadratic score $s_i : \Delta(\Theta \times \mathcal{P}_{-i}) \times \Theta \times \mathcal{P}_{-i} \to \mathbb{R}$,

$$s_i(q, \theta, E) = 2q(\theta, E) - \|q\|^2 \qquad \forall (q, \theta, E) \in \Delta(\Theta \times \mathcal{P}_{-i}) \times \Theta \times \mathcal{P}_{-i},$$

which can be readily shown to satisfy

$$s_i(q, q) - s_i(q', q) = \|q' - q\|^2 \qquad q, q' \in \Delta(\Theta \times \mathcal{P}_{-i}). \tag{24}$$

Then, define $g_i : A_i \times A_{-i} \times \Theta \to \mathbb{R}$,

$$g_i(a_i, a_{-i}, \theta) = s_i(a_i, \theta, \varphi_{-i}(a_{-i})) \qquad \forall (a_i, a_{-i}, \theta) \in A_i \times A_{-i} \times \Theta,$$

which clearly satisfies (ii).

Given $\pi_{-i} \in \Delta(A_{-i} \times \Theta)$ and $a_i \in A_i$ with $\|a_i - \pi^{\mathcal{P}}_{-i}\| < \sqrt{\varepsilon}$, for any $a'_i \in A_i$ we have

$$\begin{aligned}
g_i(a_i,&\pi_{-i}) - g_i(a'_i, \pi_{-i}) \\
&= s_i(a_i, \pi^{\mathcal{P}}_{-i}) - s_i(a'_i, \pi^{\mathcal{P}}_{-i}) && \text{(by (ii))} \\
&= s_i(a_i, \pi^{\mathcal{P}}_{-i}) - s_i(\pi^{\mathcal{P}}_{-i}, \pi^{\mathcal{P}}_{-i}) + s_i(\pi^{\mathcal{P}}_{-i}, \pi^{\mathcal{P}}_{-i}) - s_i(a'_i, \pi^{\mathcal{P}}_{-i}) \\
&= -\|a_i - \pi^{\mathcal{P}}_{-i}\|^2 + \|a'_i - \pi^{\mathcal{P}}_{-i}\|^2 && \text{(by (24))} \\
&> -\varepsilon, && \text{(by } \|a_i - \pi^{\mathcal{P}}_{-i}\| < \sqrt{\varepsilon})
\end{aligned}$$

hence $a_i \in BR_i(\pi_{-i}, G, \varepsilon)$, and this proves (iii).

Turning to (iv), let π_{-i}, $\pi'_{-i} \in \Delta(A_{-i} \times \Theta)$ with $\|\pi'^{\mathcal{P}}_{-i} - \pi^{\mathcal{P}}_{-i}\| > \zeta$ and $a_i \in A_i$ with $\|a_i - \pi^{\mathcal{P}}_{-i}\| < \sqrt{\varepsilon}$. Then, $\|a_i - \pi'^{\mathcal{P}}_{-i}\| > \zeta - \sqrt{\varepsilon} > \sqrt{3\varepsilon}$. By (i) we can find some $a'_i \in A_i$ with $\|a'_i - \pi'^{\mathcal{P}}_{-i}\| < \sqrt{\varepsilon}$. Thus,

$$\begin{aligned}
g_i(a_i,&\pi'_{-i}) - g_i(a'_i, \pi'_{-i}) \\
&= s_i(a_i, \pi'^{\mathcal{P}}_{-i}) - s_i(a'_i, \pi'^{\mathcal{P}}_{-i}) && \text{(by (ii))} \\
&= s_i(a_i, \pi'^{\mathcal{P}}_{-i}) - s_i(\pi'^{\mathcal{P}}_{-i}, \pi'^{\mathcal{P}}_{-i}) + s_i(\pi'^{\mathcal{P}}_{-i}, \pi'^{\mathcal{P}}_{-i}) - s_i(a'_i, \pi'^{\mathcal{P}}_{-i}) \\
&= -\|a_i - \pi'^{\mathcal{P}}_{-i}\|^2 + \|a'_i - \pi'^{\mathcal{P}}_{-i}\|^2 && \text{(by (24))} \\
&< -3\varepsilon + \varepsilon = -2\varepsilon, && \big(\text{by } \|a_i - \pi'^{\mathcal{P}}_{-i}\| > \sqrt{3\varepsilon} \\
& && \text{and } \|a'_i - \pi'^{\mathcal{P}}_{-i}\| < \sqrt{\varepsilon}\big)
\end{aligned}$$

and hence $a_i \notin BR_i(\pi'_{-i}, G, 2\varepsilon)$, as required. ∎

Lemma 8. *Fix $\delta > 0$ and a frame $\mathcal{P}$. Let $\varepsilon > 0$ and $G = (A_i, g_i)_{i \in I}$ satisfy the properties (i)–(iv) of Lemma 7 relative to $\zeta = \delta/(\#\Theta \cdot \max_i \#\mathcal{P}_i)$. Then the following statements hold:*

(a) for every $i \in I$ and $t_i \in \mathcal{T}_i$,

$$R(t_i, G, \varepsilon) \supseteq \Big\{a_i \in A_i \cap \varphi_i^{-1}\big(\mathcal{P}_i(t_i)\big) : \big\|a_i - \mu_i(t_i)\big|_{\Theta \times \mathcal{P}_{-i}}\big\| < \sqrt{\varepsilon}\Big\};$$

(b) for every integer $k \geqslant 0$, $i \in I$ and $s_i, t_i \in \mathcal{T}_i$ with $d^k_{i,\mathcal{P}}(s_i, t_i) > \delta$,

$$R^k_i(s_i, G, 2\varepsilon) \cap \Big\{a_i \in A_i \cap \varphi_i^{-1}\big(\mathcal{P}_i(t_i)\big) : \big\|a_i - \mu_i(t_i)\big|_{\Theta \times \mathcal{P}_{-i}}\big\| < \sqrt{\varepsilon}\Big\} = \varnothing.$$

Proof. To prove (a), we will show that the pair of correspondences $\varsigma_i : \mathcal{T}_i \rightrightarrows A_i$, $i \in I$, defined by

$$\varsigma_i(t_i) = \Big\{a_i \in A_i \cap \varphi_i^{-1}\big(\mathcal{P}_i(t_i)\big) : \big\|a_i - \mu_i(t_i)\big|_{\Theta \times \mathcal{P}_{-i}}\big\| < \sqrt{\varepsilon}\Big\} \qquad \forall t_i \in \mathcal{T}_i,$$

which is nonempty-valued by (i) of Lemma 7, has the ε-best-reply property. Indeed, given any $a_i \in \varsigma_i(t_i)$ and any conjecture $\sigma_{-i} : \Theta \times \mathcal{T}_{-i} \to \Delta(A_{-i})$ with $\operatorname{supp} \sigma_{-i}(\theta, t_{-i}) \subseteq \varsigma_{-i}(t_{-i})$ for all $(\theta, t_{-i}) \in \Theta \times \mathcal{T}_{-i}$, we must have

$$\pi_{t_i,\sigma_{-i}}\Big(\theta \times \big(A_{-i} \cap \varphi_{-i}^{-1}(E)\big)\Big) = \mu_i\big(\theta \times E | t_i\big) \qquad \forall \theta \in \Theta, E \in \mathcal{P}_{-i},$$

and hence, $\|a_i - \pi^{\mathcal{P}}_{t_i,\sigma_{-i}}\| = \big\|a_i - \mu_i(t_i)\big|_{\Theta\times\mathcal{P}_{-i}}\big\| < \sqrt{\varepsilon}$, by (i) of Lemma 7. It then follows by (iii) of Lemma 7 that $a_i \in BR_i(\pi_{t_i,\sigma_{-i}}, G, \varepsilon)$. We have thus shown that the profile of correspondences $(\varsigma_i)_{i\in I}$ has the ε-best-reply property.

Turning to (b), since $d^0_{i,\mathcal{P}} \equiv 0$ the result is true for $k = 0$ (vacuously). Proceeding by induction, assume the result is true for $k \geqslant 0$ and let us show that it remains true for $k+1$. Fix $i \in I$ and $s_i, t_i \in \mathcal{T}_i$ with $d^{k+1}_{i,\mathcal{P}}(s_i, t_i) > \delta$. Then, there is some $\theta \in \Theta$ and some $E \in \mathcal{P}_{-i}$ such that

$$\mu_i\big(\theta \times E^{\delta,k}_{\mathcal{P}}\big|s_i\big) < \mu_i\big(\theta \times E\big|t_i\big) - \delta/(\#\Theta\cdot\#\mathcal{P}_{-i}). \tag{25}$$

Consider an arbitrary k-order 2ε-rationalizable conjecture $\sigma_{-i} : \Theta\times\mathcal{T}_{-i} \to \Delta(A_{-i})$. By the induction hypothesis, for every $s_{-i}, t_{-i} \in \mathcal{T}_{-i}$, every $\theta \in \Theta$ and every $a_{-i} \in A_{-i}$ with $\big\|a_{-i} - \mu_{-i}(t_{-i})\big|_{\Theta\times\mathcal{P}_i}\big\| < \sqrt{\varepsilon}$, we can have $\sigma_{-i}(\theta, s_{-i})[a_{-i}] > 0$ only if $d^k_{-i,\mathcal{P}}(s_{-i}, t_{-i}) \leqslant \delta$. In particular, for every $s_{-i}, t_{-i} \in \mathcal{T}_{-i}$, every $\theta \in \Theta$ and every $a_{-i} \in A_{-i}$,

$$\Big(\ \mu_{-i}(t_{-i})\big|_{\Theta\times\mathcal{P}_i} = a_{-i} \quad\text{and}\quad d^k_{-i,\mathcal{P}}(s_{-i}, t_{-i}) > \delta\ \Big) \quad\Longrightarrow\quad \sigma_{-i}(\theta, s_{-i})[a_{-i}] = 0.$$

It follows that for every $(\theta, s_{-i}) \in \Theta\times\mathcal{T}_{-i}$,

$$s_{-i} \notin E^{\delta,k}_{\mathcal{P}} \quad\Longrightarrow\quad \sigma_{-i}(\theta, s_{-i})\big[A_{-i}\cap\varphi^{-1}_{-i}(E)\big] = 0, \tag{26}$$

since for every $a_{-i} \in A_{-i}\cap\varphi^{-1}_{-i}(E)$ and every $t_{-i} \in \mathcal{T}_{-i}$ with $\mu_{-i}(t_{-i})\big|_{\Theta\times\mathcal{P}_i} = a_{-i}$ we must have $t_{-i} \in E$. Hence, for every k-order 2ε-rationalizable conjecture σ_{-i} we have

$$\begin{aligned}
&\pi_{s_i,\sigma_{-i}}\Big(\theta\times\big(A_{-i}\cap\varphi^{-1}_{-i}(E)\big)\Big)\\
&\qquad= \int_{\mathcal{T}_{-i}} \sigma_{-i}(\theta, s_{-i})\big[A_{-i}\cap\varphi^{-1}_{-i}(E)\big]\,\mu_i\big(\theta\times ds_{-i}\big|s_i\big)\\
&\qquad= \int_{E^{\delta,k}_{\mathcal{P}}} \sigma_{-i}(\theta, s_{-i})\big[A_{-i}\cap\varphi^{-1}_{-i}(E)\big]\,\mu_i\big(\theta\times ds_{-i}\big|s_i\big) && \text{(by (26))}\\
&\qquad\leqslant \mu_i\big(\theta\times E^{\delta,k}_{\mathcal{P}}\big|s_i\big)\\
&\qquad< \mu_i(\theta\times E|t_i) - \delta/(\#\Theta\cdot\#\mathcal{P}_{-i}), && \text{(by (25))}
\end{aligned}$$

and thus, $\big\|\pi^{\mathcal{P}}_{s_i,\sigma_{-i}} - \mu_i(t_i)\big|_{\Theta\times\mathcal{P}_{-i}}\big\| > \zeta$. It follows from part (iv) of Lemma 7 that

$$R^{k+1}_i(s_i, G, 2\varepsilon)\cap\Big\{a_i : \big\|a_i - \mu_i(t_i)\big|_{\Theta\times\mathcal{P}_{-i}}\big\| < \sqrt{\varepsilon}\Big\} = \varnothing,$$

as was to be shown. ■

Corollary 2. *For every frame $\mathcal{P}$ and $\delta > 0$ there exists $\varepsilon > 0$ and a game G with $|G| \leqslant M$ such that for every $i \in I$ and $s_i, t_i \in \mathcal{T}_i$,*

$$d^{UW}_{i,\mathcal{P}}(s_i, t_i) > \delta \quad\Longrightarrow\quad R_i(t_i, G) \not\subseteq R_i(s_i, G, \varepsilon).$$

Proof. A straightforward implication of Lemma 8 is that for some $\varepsilon > 0$ and some game G' with $|G'| \leqslant M$, for every $i \in I$ and $s_i, t_i \in \mathcal{T}_i$ we have

$$d^{UW}_{i,\mathcal{P}}(s_i, t_i) > \delta \quad \Longrightarrow \quad R_i(t_i, G', \varepsilon) \not\subseteq R_i(s_i, G', 2\varepsilon).$$

To conclude, we can use a construction similar to the last part of the proof of Lemma 4 to obtain a game G with $R_i(\cdot, G, \gamma) = R_i(\cdot, G', \gamma + \varepsilon)$ for all $\gamma \geqslant 0$. ∎

The only if direction of Theorem 2 is an immediate implication of Corollary 2.

A.4 Proof of Theorem 3

We need the following piece of notation. Given an integer $m \geqslant 1$, a measurable subset $E \subseteq \Omega$ and $\mathbf{p} \in [0,1]^2$, define the *event that E is m-order $\mathbf{p}$-belief* recursively as follows:

$$\left[B^{\mathbf{p}}\right]^m(E) \overset{\text{def}}{=} B^{\mathbf{p}}\left(E \cap \left[B^{\mathbf{p}}\right]^{m-1}(E)\right),$$

where $\left[B^{\mathbf{p}}\right]^0(E) \overset{\text{def}}{=} \Omega$. Then, the *event that E is m-order $\mathbf{p}$-belief for player i*, written $\left[B^{\mathbf{p}}_i\right]^m(E)$, is the projection of $\left[B^{\mathbf{p}}\right]^m(E)$ onto $\mathcal{T}_i$. We these definitions in place, we have:

$$C^{\mathbf{p}}(E) = \bigcap_{m \geqslant 1} \left[B^{\mathbf{p}}\right]^m(E) \quad \text{and} \quad C^{\mathbf{p}}_i(E) = \bigcap_{m \geqslant 1} \left[B^{\mathbf{p}}_i\right]^m(E).$$

Lemma 9. *For every integer $k \geqslant 1$ and $\delta > 0$ there exists a k-order frame $\mathcal{P}$ such that, for every $i \in I$, every atom of $\mathcal{P}_i$ has d^k_i-diameter at most δ.*

Proof. For $k = 1$ the result is trivial, as any profile of first-order measurable partitions is a frame. Proceeding by induction, consider $k \geqslant 1$, fix $\delta > 0$ and let $\mathcal{P}$ be a k-order measurable frame whose atoms all have d^k_i-diameter less than $\delta/2$. Let Π_i be a finite partition of the simplex $\Delta(\Theta \times \mathcal{P}_{-i})$ (viewed as a subset of the Euclidean space $\mathbb{R}^{\#\Theta \cdot \#\mathcal{P}_{-i}}$) into finitely many Borel measurable subsets with Euclidean diameter less than $\delta/\sqrt{\#\Theta\#\mathcal{P}_{-i}}$. By Lemma 2, the join $(\mathcal{P}_i \vee (\Pi_i/\mathcal{T}_i))_{i \in I}$ is a $(k+1)$-order frame. We claim that every atom of $\mathcal{P}_i \vee (\Pi_i/\mathcal{T}_i)$ has d^{k+1}_i-diameter less than δ. Let s_i, t_i be a pair of types with $s_i \in \left(\mathcal{T}_i/\Pi_i\right)(t_i)$, let E be a measurable subset of $\Theta \times \mathcal{T}_{-i}$ and let us show that $\mu_i(E^{\delta,k}|s_i) \geqslant \mu_i(E|t_i) - \delta$. Since all the atoms of $\mathcal{P}_{-i}$ have d^k_i-diameter less than $\delta/2$, there is some $F \in 2^\Theta \otimes \mathcal{P}_{-i}$ with $E \subseteq F \subseteq E^{\delta/2,k}$. Then, $|\mu_i(F|s_i) - \mu_i(F|t_i)| < \delta$, because every atom of Π_i has Euclidean diameter less than $\delta/\sqrt{\#\Theta\#\mathcal{P}_{-i}}$ and we have $s_i \in \left(\mathcal{T}_i/\Pi_i\right)(t_i)$. Thus,

$$\begin{aligned}
\mu_i(E^{\delta,k}|s_i) &\geqslant \mu_i(F^{\delta/2,k}|s_i) && (\text{by } F \subseteq E^{\delta/2,k}) \\
&\geqslant \mu_i(F|s_i) && (\text{by } F \subseteq F^{\delta/2,k}) \\
&\geqslant \mu_i(F|t_i) - \delta && (\text{by } |\mu_i(F|s_i) - \mu_i(F|t_i)| < \delta) \\
&\geqslant \mu_i(E|t_i) - \delta && (\text{by } E \subseteq F)
\end{aligned}$$

as claimed. ∎

Lemma 10. *Fix $\delta > 0$, an integer $k \geqslant 1$ and a k-order frame $\mathcal{P}$ whose atoms all have d_i^k-diameter less than δ for every $i \in I$. Then, for every $m = 0, \ldots, k$,*

$$d_i^m(s_i, t_i) \leqslant d_{i,\mathcal{P}}^m(s_i, t_i) + m\delta \qquad \forall s_i, t_i \in \mathcal{T}_i, \forall i \in I.$$

Proof. Fix $\delta > 0$, an integer $k \geqslant 1$ and a k-order frame $\mathcal{P}$ whose atoms all have d_i^k-diameter less than δ, for every $i \in I$. (Such a frame exists by Lemma 9.) For $m = 0$ the conclusion of the lemma is trivial, as $d_i^0 = d_{i,\mathcal{P}}^0 = 0$. Consider $1 \leqslant m \leqslant k$ and assume the conclusion of the lemma holds for $m-1$. Let $s_i, t_i \in \mathcal{T}_i$ and $\eta > d_{i,\mathcal{P}_i}^m(s_i, t_i)$ and let us show that $d_i^m(s_i, t_i) \leqslant \eta + m\delta$. Fix $E \in 2^\Theta \otimes \mathcal{P}_{-i}$. Since all the atoms of $\mathcal{P}_{-i}$ have d_{-i}^k-diameter less than δ, there is some $F \in 2^\Theta \otimes \mathcal{P}_{-i}$ with $E \subseteq F \subseteq E^{\delta, m-1}$. Then,

$$\begin{aligned}
\mu_i(E^{\eta+m\delta,m-1}|s_i) &\geqslant \mu_i(F^{\eta+(m-1)\delta,m-1}|s_i) && \text{(by } F \subseteq E^{\delta,m-1}\text{)} \\
&\geqslant \mu_i(F_{\mathcal{P}}^{\eta,m-1}|s_i) && \text{(by } d_{-i}^{m-1} \leqslant d_{-i,\mathcal{P}}^{m-1} + (m-1)\delta\text{)} \\
&\geqslant \mu_i(F|t_i) - \eta && \text{(by } \eta > d_{i,\mathcal{P}_i}^m(s_i,t_i)\text{)} \\
&\geqslant \mu_i(E|t_i) - \eta - m\delta, && \text{(by } E \subseteq F\text{)}
\end{aligned}$$

and hence, $d_i^m(s_i, t_i) \leqslant \eta + m\delta$. But since our choice of $\eta > d_{i,\mathcal{P}_i}^m(s_i, t_i)$ was arbitrary, we have shown that $d_i^m(s_i, t_i) \leqslant d_{i,\mathcal{P}_i}^m(s_i, t_i) + m\delta$, as required. ∎

Lemma 11. *Fix $\delta > 0$, $\mathbf{p} \in (0, 1]^2$, an integer $k \geqslant 1$ and an event $E \subseteq \Omega$ that is measurable with respect to a frame $\tilde{\mathcal{P}}$. Let $\mathcal{P}$ denote the common belief frame $\left(\tilde{\mathcal{P}}_i \vee \left\{C_i^{\mathbf{p}}(E), \mathcal{T}_i \setminus C_i^{\mathbf{p}}(E)\right\}\right)_{i \in I}$. Then, for every integer $\ell \geqslant k$,*

$$\left(C_i^{\mathbf{p}}(E)\right)_{\mathcal{P}}^{\delta,\ell} \subseteq \left[B_i^{\mathbf{p}-\delta\mathbf{1}}\right]^{\ell-k}\left(E_{\tilde{\mathcal{P}}}^{\delta,k}\right) \qquad \forall i \in I.$$

Proof. We prove the result by induction on $\ell \geqslant k$. First, the result is trivial for $\ell = k$, as $\left[B_i^{\mathbf{p}-\delta\mathbf{1}}\right]^{\mathbf{0}}(\cdot) \equiv \mathcal{T}_i$ for every i. Next, suppose the result is true for $\ell \geqslant k$ and let us show that it remains true for $\ell+1$. Pick $t_i \in C_i^{\mathbf{p}}(E)$ and $s_i \in \mathcal{T}_i$ with $d_{i,\mathcal{P}}^{\ell+1}(s_i, t_i) < \delta$ and let us show that

$$s_i \in B_i^{p_i-\delta}\left(E_{\tilde{\mathcal{P}}}^{\delta,k} \cap \left[B^{\mathbf{p}-\delta\mathbf{1}}\right]^{\ell-k}\left(E_{\tilde{\mathcal{P}}}^{\delta,k}\right)\right) = \left[B_i^{\mathbf{p}-\delta\mathbf{1}}\right]^{\ell+1-k}\left(E_{\tilde{\mathcal{P}}}^{\delta,k}\right). \tag{27}$$

Indeed,[34]

$$\begin{aligned}
\mu_i\big(\big(E^{\delta,k}_{\tilde{\mathcal{P}}} \cap \big[B^{\mathbf{p}-\delta\mathbf{1}}\big]^{\ell-k}\big(E^{\delta,k}_{\tilde{\mathcal{P}}}\big)\big)_{s_i}\big|s_i\big) &\geqslant \mu_i\big(\big(\big(E \cap C^{\mathbf{p}}(E)\big)^{\delta,\ell}_{\mathcal{P}}\big)_{s_i}\big|s_i\big) \\
&\geqslant \mu_i\big(\big(\big(E \cap C^{\mathbf{p}}(E)\big)_{t_i}\big)^{\delta,\ell}_{\mathcal{P}}\big|s_i\big) \\
&\geqslant \mu_i\big(\big(E \cap C^{\mathbf{p}}(E)\big)_{t_i}\big|t_i\big) - \delta \\
&\geqslant \mu_i\big(\big(E \cap C^{\mathbf{p}}(E)\big)_{t_i}\big|t_i\big) - \delta \\
&\geqslant p_i - \delta,
\end{aligned}$$

where the first inequality follows from the induction hypothesis and the fact that $d^k_{i,\tilde{\mathcal{P}}} \leqslant d^\ell_{i,\mathcal{P}}$, the second inequality follows from $d^\ell_{i,\mathcal{P}}(s_i, t_i) < \delta$, the third inequality follows from $d^{\ell+1}_{i,\mathcal{P}}(s_i, t_i) < \delta$ and the last inequality follows from $t_i \in C^{\mathbf{p}}_i(E) = B^{p_i}_i(E \cap C^{\mathbf{p}}(E))$. This proves (27). ■

Lemma 12. *For every $\delta > 0$, integer $k \geqslant 1$, $\mathbf{p} \in (0,1]^2$ and measurable $E \subseteq \Omega$ there exists a frame $\mathcal{P}$ such that for every $s_i, t_i \in \mathcal{T}_i$ with $d^{UW}_{i,\mathcal{P}}(s_i, t_i) < \delta$,*

$$t_i \in C^{\mathbf{p}}_i(E) \quad\Longrightarrow\quad s_i \in C^{\mathbf{p}-\delta(k+2)\mathbf{1}}_i\big(E^{\delta(k+2),k}\big).$$

Proof. Fix $\varepsilon > 0$, an integer $k \geqslant 1$, $\mathbf{p} \in (0,1]^2$ and a measurable set $E \subseteq \Omega$. Let $\tilde{\mathcal{P}}$ be a k-order frame whose atoms have d^k_i-diameter at most δ, for every player i. Fix $i \in I$, pick $F \in 2^\Theta \otimes \mathcal{P}_{-i}$ such that $E \subseteq F \subseteq E^{\delta,k}$ and consider the common belief frame $\mathcal{P} = \big(\tilde{\mathcal{P}}_j \vee \big\{C^{\mathbf{p}}_j(F), \mathcal{T}_j \setminus C^{\mathbf{p}}_j(F)\big\}\big)_{j\in I}$. Let $t_i \in C^{\mathbf{p}}_i(E)$ and consider a type s_i with $d^{UW}_{i,\mathcal{P}}(s_i, t_i) < \delta$. Since $E \subseteq F$, we have $t_i \in C^{\mathbf{p}}_i(F)$. Then,

$$\begin{aligned}
s_i \in \bigcap_{\ell \geqslant k} \big(C^{\mathbf{p}}_i(F)\big)^{\delta,\ell}_{\mathcal{P}} &\subseteq \bigcap_{m \geqslant 1} \Big[B^{\mathbf{p}-\delta\mathbf{1}}_i\Big]^m \big(F^{\delta,k}_{\tilde{\mathcal{P}}}\big) && \text{(by } d^{UW}_{i,\mathcal{P}}(s_i,t_i) < \delta \text{ and Lemma 11)} \\
&= C^{\mathbf{p}-\delta\mathbf{1}}_i\big(F^{\delta,k}_{\tilde{\mathcal{P}}}\big) && \text{(by the definition of common belief)} \\
&\subseteq C^{\mathbf{p}-\delta\mathbf{1}}_i\big(F^{(k+1)\delta,k}\big) && \text{(by Lemma 10)} \\
&\subseteq C^{\mathbf{p}-\delta\mathbf{1}}_i\big(E^{(k+2)\delta,k}\big), && \text{(by } F \subseteq E^{\delta,k}\text{)}
\end{aligned}$$

and hence $s_i \in C^{\mathbf{p}-(k+2)\delta\mathbf{1}}_i\big(E^{(k+2)\delta,k}\big)$. ■

Theorem 3 is an immediate implication of Lemma 12 above.

[34] Recall that, for any measurable subset $E \subseteq \Omega$ and any type t_i of player i, E_{t_i} denotes the section of E over t_i. See Section 2.3.

A.5 Proof of Claim 1

First we introduce some notation. For every $\mathbf{r} = (r_1, \ldots, r_4) \in (0,1]^4$, $i \in I$, measurable subset $E \subseteq \Omega$, let $\left[M_i^{\mathbf{r}}\right]^0(E) = \Omega$ and for every positive integer k, define recursively

$$\left[M_i^{\mathbf{r}}\right]^k(E) = B_i^{r_1}\big(E \cap B_{-i}^{r_2}\big(E \cap B_i^{r_3}\big(E \cap B_{-i}^{r_4}\big(E \cap \left[M_i^{\mathbf{r}}\right]^{k-1}(E)\big)\big)\big)\big).$$

Note that

$$\left[M_i^{\mathbf{r}}\right]^k(E) \;\subseteq\; B_i^{r_1}\Big(E \cap \Big[M_{-i}^{(r_2,r_3,r_4,r_1)}\Big]^{k-1}(E)\Big) \;\subseteq\; \left[M_i^{\mathbf{r}}\right]^{k-1}(E) \qquad \forall k \geqslant 1. \tag{28}$$

Next, we define the frame $\mathcal{P}$. For every $k \geqslant 1$, let

$$P_1^k = \Big[M_1^{(q,p,p,p)}\Big]^k(\theta_0), \qquad Q_1^k = \Big[M_1^{(p,p,q,p)}\Big]^k(\theta_0),$$
$$P_2^k = \Big[M_2^{(p,p,p,q)}\Big]^k(\theta_0), \qquad Q_2^m = \Big[M_2^{(p,q,p,p)}\Big]^k(\theta_0)$$

and finally,

$$P_i = \bigcap_{k=1}^{\infty} P_i^k \qquad \text{and} \qquad Q_i = \bigcap_{k=1}^{\infty} Q_i^k \qquad \forall i \in I.$$

The frame $\mathcal{P}$ is the profile of partitions $(\mathcal{P}_i)_{i \in I}$ where

$$\mathcal{P}_i \stackrel{\text{def}}{=} \big\{P_i \setminus Q_i,\, Q_i \setminus P_i,\, P_i \cap Q_i,\, \mathcal{T}_i \setminus (P_i \cup Q_i)\big\} \qquad \forall i \in I.$$

To verify that $\mathcal{P}$ is a frame, note that (28) implies

$$P_1^k \;\subseteq\; B_1^q(\theta_0 \times P_2^{k-1}) \;\subseteq\; P_1^{k-1} \qquad \forall k \geqslant 1$$

and hence

$$P_1 \;=\; \bigcap_{k=2}^{\infty} P_1^k \;\subseteq\; \bigcap_{k=2}^{\infty} B_1^q(\theta_0 \times P_2^{k-1}) \;=\; B_1^q(\theta_0 \times P_2) \;\subseteq\; \bigcap_{k=2}^{\infty} P_1^{k-1} \;=\; P_1.$$

Therefore, $P_1 = B_1^q(\theta_0 \times P_2)$. By analogous arguments,

$$Q_1 = B_1^p(\theta_0 \times Q_2), \qquad P_2 = B_2^p(\theta_0 \times Q_1), \qquad \text{and} \qquad Q_2 = B_2^p(\theta_0 \times P_1).$$

Thus, for all $i \in I$ and $t_i \in \mathcal{T}_i$, the element of $\mathcal{P}_i$ containing t_i is determined by the values $\mu_i(\theta_0 \times P_{-i}|t_i)$ and $\mu_i(\theta_0 \times Q_{-i}|t_i)$, and hence, a fortiori, by the restriction of $\mu_i(\cdot|t_i)$ to $\Theta \times \mathcal{P}_{-i}$.

We now prove that $t_1^n \not\to t_1$ uniform-weakly on $\mathcal{P}$. Fix $0 < \delta < \min\{p, \Delta\}$. It is enough to show that for every positive integer n,

$$d_{1,\mathcal{P}}^{4n-3}(t_1^n, t_1) > \delta, \quad d_{2,\mathcal{P}}^{4n-2}(s_2^n, s_2) > \delta, \quad d_{1,\mathcal{P}}^{4n-1}(s_1^n, s_1) > \delta, \quad d_{2,\mathcal{P}}^{4n}(t_2^n, t_2) > \delta. \tag{29}$$

To prove this, we will show that for $n = 1$, the first inequality holds, and that for any $n \geqslant 1$, if the first inequality holds for n, then all others also hold for n, and if the last inequality holds for

n, then the first holds for $n+1$. In the proof we will use the following facts, which are immediate from the definition of t_1, t_2, s_1, s_2 and the fact that $\mu_i(\theta_0|r_i) = 0$ and $P_i \cup Q_i \subseteq B_i^p(\theta_0)$ for each $i \in I$:

$$t_1 \in P_1, \quad s_1 \in Q_1, \quad s_2 \in Q_2, \quad t_2 \in P_2, \quad r_i \notin (P_i \cup Q_i)_{\mathcal{P}}^{\delta,1} \quad \forall i \in I. \tag{30}$$

Since $\mu_1(\theta_0|t_1) = 1$ and $\mu_1(\theta_0|t_1^1) = 0$, we have $d_{1,\mathcal{P}}^1(t_1^1, t_1) = 1 > \delta$, which proves the first inequality in (29) for $n = 1$. Now fix any $n \geqslant 1$ and assume the first inequality in (29) holds for n. Then, by (30), $\{t_1^n, r_1\} \cap (P_1)_{\mathcal{P}}^{\delta,4n-3} = \varnothing$. Since $\mu_2(\theta_0 \times \{t_1^n, r_1\}|s_1^n) = 1$ and, by (30), $\mu_2(\theta_0 \times P_1|s_2) = \mu_2(\theta_0, t_1|s_2) = p$, it follows that $d_{2,\mathcal{P}}^{4n-2}(s_2^n, s_2) \geqslant p > \delta$. Thus, the second inequality in (29) holds for n. The latter implies, by (30), $\{s_2^n, r_2\} \cap (P_2)_{\mathcal{P}}^{\delta,4n-2} = \varnothing$, and since $\mu_1(\theta_0 \times \{s_2^n, r_2\}|s_1^n) = 1$ and, by (30), $\mu_1(\theta_0 \times Q_2|s_1) = \mu_1(\theta_0, s_2|s_1) = p$, we also have $d_{1,\mathcal{P}}^{4n-1}(s_1^n, s_1) \geqslant p > \delta$, i.e., the third inequality in (29) holds for n. This in turn implies, by (30), $\{s_1^n, r_1\} \cap (P_1)_{\mathcal{P}}^{\delta,4n-1} = \varnothing$, and since $\mu_2(\theta_0 \times \{s_1^n, r_1\}|t_2^n) = 1$ and, by (30), $\mu_2(\theta_0 \times Q_1|t_2) = \mu_1(\theta_0, s_1|t_2) = p$, we obtain $d_{2,\mathcal{P}}^{4n}(t_2^n, s_2) \geqslant p > \delta$. This proves that the fourth inequality in (29) holds for n, and hence, by (30), that $\{t_2^n, r_2\} \cap (P_2)_{\mathcal{P}}^{\delta,4n} = \varnothing$. It remains to show that the latter implies that the first inequality in (29) holds for $n+1$. Indeed, since $\mu_1(\theta_0 \times \{t_2^n, r_2\}|t_1^{n+1}) = 1 - q + \Delta$ and, by (30), $\mu_1(\theta_0 \times P_2|t_1) = \mu_1(\theta_0, t_2|t_1) = q$, we have $d_{1,\mathcal{P}}^{4n+1}(t_1^{n+1}, t_1) \geqslant \Delta > \delta$, as claimed.

A.6 Proof of Theorem 4

It is enough to prove the implication $(e) \Rightarrow (a)$, since the implications $(a) \Leftrightarrow (b) \Rightarrow (c) \Leftrightarrow (d) \Rightarrow (e)$ follow from Theorems 1, 2 and 3. Fix a finite type space (T_1, T_2). We will show that there exist $\eta > 0$ and $k \geqslant 1$ such that for all $\delta \in (0, \eta)$, $n \geqslant 0$, $i \in I$ and $(t_i, s_i) \in T_i \times \mathcal{T}_i$ with $d_i^k(s_i, t_i) < \delta$, one has

$$s_i \in \left[B_i^{(1-\delta)\mathbf{1}}\right]^n \left(\Theta \times T_1^{\delta,k} \times T_2^{\delta,k}\right) \quad \Rightarrow \quad d_i^{k+n}(s_i, t_i) < 2\delta\,. \tag{31}$$

Choose $k \geqslant 1$ and $\eta > 0$ so that, for all $i \in I$ and $s_i, t_i \in T_i$, if $d_i^{k-1}(s_i, t_i) < 2\eta$ then $s_i = t_i$. Since T_1 and T_2 are finite, such k and η exist. Fix $\delta \in (0, \eta)$. Thus, for all $n \geqslant 0$, $i \in I$, $t_i \in T_i$ and $E \subseteq \Theta \times T_{-i}$,

$$E^{\delta,k-1} \cap (\Theta \times T_{-i})^{\delta,k} \subseteq E^{\delta,k}. \tag{32}$$

The proof of (31) is by induction in n. Obviously, (31) holds for $n = 0$. Now assume that (31) holds for some $n \geqslant 0$, and fix $i \in I$, $t_i \in T_i$ and $s_i \in \mathcal{T}_i$ with $d_i^k(s_i, t_i) < \delta$ and $s_i \in \left[B_i^{(1-\delta)\mathbf{1}}\right]^{n+1} \left(\Theta \times T_1^{\delta,k} \times T_2^{\delta,k}\right)$. Since $\mu_i(\Theta \times T_{-i}|t_i) = 1$, in order to prove that $d_i^{k+n+1}(s_i, t_i) < 2\delta$ it suffices to show that $\mu_i(E^{2\delta,k+n}|s_i) \geqslant \mu_i(E|t_i) - 2\delta$ for every $E \subseteq$

$\Theta \times T_{-i}$. Indeed,

$$\begin{aligned}
\mu_i\left(E^{2\delta,k+n} \middle| s_i\right) &\geq \mu_i\left(E^{\delta,k} \cap \left[B_{-i}^{(1-\delta)\mathbf{1}}\right]^n \left(\Theta \times T_1^{\delta,k} \times T_2^{\delta,k}\right) \middle| s_i\right) \\
&\geq \mu_i\left(E^{\delta,k-1} \cap (\Theta \times T_{-i})^{\delta,k} \cap \left[B_{-i}^{(1-\delta)\mathbf{1}}\right]^n \left(\Theta \times T_1^{\delta,k} \times T_2^{\delta,k}\right) \middle| s_i\right) \\
&\geq \mu_i\left(E^{\delta,k-1} \middle| s_i\right) - \delta \geq \mu_i\left(E \middle| t_i\right) - 2\delta,
\end{aligned}$$

where the first inequality follows from the induction hypothesis, the second from (32), the third from $s_i \in \left[B_i^{(1-\delta)\mathbf{1}}\right]^{n+1}(T^{\delta,k})$ and the fourth from $d_i^k(s_i, t_i) < \delta$.

A.7 Proof of Theorem 5

An immediate implication of Ely and Pęski (2011, Theorem 1) is that every finite type is critical. This fact, together with the denseness of finite types in the strategic topology (proved in Dekel, Fudenberg, and Morris (2006)), implies that the set of critical types is dense in the strategic topology.

Next, we show that the set of critical types is open in the strategic topology, or equivalently, that the set of regular types is closed. Suppose not. Then, there is a sequence of regular types t_i^n that converges to some critical type t_i. By Ely and Pęski (2011, Theorem 1), there is some $p > 0$ and some product-closed, proper subset $E \subset \Omega$ with $t_i \in C_i^{p\mathbf{1}}(E)$. Then, there is an integer $k \geqslant 1$ and $\delta \in (0, p)$ such that the d_i^k-closure of $E^{\delta,k}$ is a proper subset of Ω. Moreover, by Theorem 3, $t_i^n \to t_i$ in common beliefs, and hence $t_i^n \in C_i^{(p-\delta)\mathbf{1}}(E^{\delta,k})$ for all n large enough. It follows, again by Ely and Pęski (2011, Theorem 1), that t_i^n is a critical type for all n large enough, and this is a contradiction. The contradiction shows that the set of regular types is closed.

Finally, to prove that the set of critical types is finitely prevalent, given a pair of types u_i, u_i' and $\alpha \in (0, 1)$, define

$$u_i'' = \alpha u_i + (1-\alpha) u_i' = \mu_i^{-1}\left(\alpha \mu_i(u_i) + (1-\alpha)\mu_i(u_i')\right).$$

We show that u_i'' is regular if and only if u_i and u_i' are both regular. To prove the "only if" part, suppose that u_i is critical. By Ely and Pęski (2011, Theorem 1), $u_i \in C_i^{p\mathbf{1}}(E_i) = E_i \cap B_i^p(C_{-i}^{p\mathbf{1}}(E_i))$ for some product-closed proper subset $E_i \subset \mathcal{T}_i$, and this clearly implies $u_i'' \in (E_i \cup u_i'') \times B_i^{\alpha p}(C_{-i}^{\alpha p\mathbf{1}}(E_i \cup u_i'')) = C_i^{\alpha p\mathbf{1}}(E_i \cup u_i'')$, which implies u_i'' is critical. Conversely, if u_i'' is critical, then, by Theorem 1 in Ely and Pęski (2011), $u_i'' \in C_i^p(E_i) = E_i \cap B_i^p(C_{-i}^{p\mathbf{1}})$ for some product-closed proper subset $E_i \subset \mathcal{T}_i$ and some $p > 0$. Since $u'' = \alpha u_i + (1-\alpha)u_i'$, either $u_i \in B_i^p(C_{-i}^{p\mathbf{1}}(E_i))$ or $u_i' \in B_i^p(C_{-i}^{p\mathbf{1}}(E_i))$, thus either u_i or u_i' is critical. The conclusion that the set of critical types is finitely prevalent now follows from Heifetz and Neeman (2006, Lemma 1).

A.8 Proof of Theorem 6

The proof of the theorem uses the following lemma, a generalization of Theorem A in Monderer and Samet (1989), which in turn generalizes Aumann's (1976) agreement theorem. It says that if a common prior type commonly (p, p)-believes, for some $p > 0$, that each player i's belief in a fixed event E lies in an interval of a given size δ around some fixed value r_i, then the difference between the values r_1 and r_2 is bounded above by a number that vanishes for sufficiently large p and small δ.[35]

Lemma 13. *Let $r_1, r_2 \in [0,1]$, $p \in (0,1]$ and $\delta \in [0,1]$. Let $D \subseteq \Theta \times \mathcal{T}_1 \times \mathcal{T}_2$ be a measurable set and define $E_i = \{t_i \in \mathcal{T}_i : |\mu_i(D|t_i) - r_i| \leqslant \delta\}$, $i \in I$. If, for some i, $C_i^{p\mathbf{1}}(\Theta \times E_1 \times E_2)$ contains a common prior type of player i, then $|r_1 - r_2| \leqslant 1 - p + 2\delta$.*

Proof. Let $F = C^{p\mathbf{1}}(\Theta \times E_1 \times E_2)$ and $F_i = C_i^{p\mathbf{1}}(\Theta \times E_1 \times E_2)$ for each player i. Suppose $F_i \cap T_i \neq \emptyset$ for some i, where $(T_j)_{j\in I}$ is a countable common prior type space with prior η. Then, since $p > 0$, we must have $F_i \cap T_i \neq \emptyset$ for every player i. Moreover, we have $F_i \subseteq E_i$. Thus, since $\eta(\cdot|t_i) = \mu_i(\cdot|t_i)$ for every $t_i \in T_i$, it follows that $|\eta(D|F_i) - r_i| \leqslant \delta$. Similarly, since $F_i \subseteq B_i^p(F)$, we have $\eta(F|F_i) \geqslant p$. Thus, since $F \subseteq F_i$, for any event $D' \subseteq \Omega$ we get

$$\eta(D'|F_i) \geqslant \eta(D' \cap F|F_i) = \eta(F|F_i)\eta(D'|F) \geqslant p\eta(D'|F).$$

Plugging first D and then its complement for D' in the latter expression, we obtain

$$r_i + \delta \geqslant \eta(D|F_i) \geqslant p\eta(D|F), \qquad r_i - \delta \leqslant \eta(D|F_i) \leqslant 1 - p + p\eta(D|F),$$

and hence $p\eta(D|F) - \delta \leqslant r_i \leqslant p\eta(D|F) + 1 - p + \delta$. Since the latter inequalities hold for every player i, the result follows. ■

As a preliminary step, we prove the existence of a finite type t_i that cannot be approximated uniform weakly by any sequence of common prior types. Let θ_1 and θ_2 be distinct elements of Θ. Consider the type space $(T_i)_{i\in I}$ where, for each $i \in I$, T_i is a singleton, $T_i = \{t_i\}$, and beliefs are such that $\mu_i(\theta_i, t_{-i}|t_i) = 1$. Fix $\delta \in (0, 1/3)$. Letting $D = \{\theta_1\} \times \mathcal{T}_1 \times \mathcal{T}_2$, $p = 1 - \delta$, $r_1 = 1$ and $r_2 = 0$ in Lemma 13, we see that for each player i, $C_i^{p\mathbf{1}}(B_1^p(\theta_1) \cap B_2^p(\theta_2))$ contains no common prior types of i. Since $t_i \in C^{\mathbf{1}}(B_1^1(\theta_1) \cap B_2^1(\theta_2))$, we conclude that that there is no sequence of common prior types converging to t_i in common beliefs, and hence, by Theorem 3, no sequence of common prior types converging to t_i uniform weakly. Thus, we can find $\varepsilon \in (0,1)$ such that $d_i^{UW}(t_i, s_i) > \varepsilon$ for every i and every common prior type s_i.

Now, to show that the set of common prior types is nowhere dense in the universal type space under the strategic topology, fix any $u_i \in \mathcal{T}_i$. We prove the existence of a sequence of finite

[35] Neeman (1996) proves the claim for $\delta = 0$, improving on the bound originally given by Monderer and Samet (1989). Aumann (1976) proves the result for $\delta = 0$ and $p = 1$.

types, t_i^n, that converges to u_i strategically and satisfies $d_i^{UW}(t_i^n, s_i) > \varepsilon/n$ for every n and every common prior type s_i. Since each t_i^n is finite, this implies, by Theorem 4, that no t_i^n is in the strategic closure of the set of common prior types, and hence that the complement of that closure is dense, thus proving the nowhere denseness of common prior types. To define the sequence of finite types t_i^n, pick an arbitrary sequence of finite types, u_i^n, converging to u_i strategically—such sequence exists by Theorem 3 in Dekel, Fudenberg, and Morris (2006)—, and for each n define t_i^n by

$$\mu_i(\cdot|t_i^n) = \frac{n-1}{n}\mu_i(\cdot|u_i^n) + \frac{1}{n}\mu_i(\cdot|t_i).$$

Since u_i^n and t_i are finite, so is t_i^n. Moreover, $d_i^{UW}(t_i^n, u_i^n) \leqslant 1/n$ for every n. Since the sequence u_i^n converges strategically to u_i, we conclude by Theorem 1 that the sequence t_i^n also converges strategically to u_i. It remains to prove that $d_i^{UW}(t_i^n, s_i) > \varepsilon/n$ for every n and every common prior type s_i. Indeed, assume, on the contrary, that $d_i^{UW}(t_i^n, s_i) \leqslant \varepsilon/n$ for some n and common prior s_i. Then, since $\mu_i(t_{-i}|t_i^n) \geqslant 1/n$, we have $\mu_i(\{t_{-i}\}^{\varepsilon/n,k}|s_i) \geqslant (1-\varepsilon)/n$ for every $k \geqslant 1$. Thus, we also have $\mu_i(\{s_{-i} \in \mathcal{T}_{-i} : d_{-i}^{UW}(s_{-i}, t_{-i}) \leqslant \varepsilon/n\}|s_i) = \mu_i(\cap_{k \geqslant 1}\{t_{-i}\}^{\varepsilon/n,k}|s_i) \geqslant (1-\varepsilon)/n > 0$. Since s_i is a common prior type of player i, $\mu_i(\cdot|s_i)$ is supported in the set of common prior types of player $-i$. Thus, $d_{-i}^{UW}(t_{-i}, s_{-i}) \leqslant \varepsilon/n \leqslant \varepsilon$ for some common prior type s_{-i}. This is a contradiction, as we have shown in the preliminary step that there is no common prior type that is ε-close to t_{-i} in the uniform weak distance.

We now show that the set of common prior types is finitely shy, thus concluding the proof of the theorem. For each player i, consider the set of types defined by

$$T_i^0 = \{u_i \in \mathcal{T}_i : \mu_i(\Theta \times \{t_{-i}\}|u_i) = 0\},$$

where $(\{t_i\})_{i\in I}$ is the singleton type space defined in the preliminary step above. Next, for all $u_i, u_i' \in \mathcal{T}_i$ and $\alpha \in (0,1)$, define

$$u_i'' = \alpha u_i + (1-\alpha)u_i' = \mu_i^{-1}\big(\alpha\mu_i(\cdot|u_i) + (1-\alpha)\mu_i(\cdot|u_i')\big).$$

The set T_i^0 is a proper face of $\mathcal{T}_i$, i.e. $\alpha u_i + (1-\alpha)u_i' \in T_i^0$ if and only if $u_i, u_i' \in T_i^0$. Thus, by Heifetz and Neeman (2006, Lemma 1), the set T_i^0 is finitely shy.

We now prove that the set of common prior types is contained in T_i^0 by showing that for every type u_i the inequality $\mu_i(\Theta \times \{t_{-i}\}|u_i) > 0$ implies that u_i is not a common prior type. Suppose that u_i is a common prior type contained in a countable common prior type space $(T_i')_{i\in I}$ with prior η. Since $\eta_i(u_i) > 0$ and $\mu_i(\Theta \times \{t_{-i}\}|u_i) > 0$, it follows that $t_i \in T_i'$ and $t_{-i} \in T_{-i}'$. Thus, $\eta_i(t_i) > 0$ and $\eta_{-i}(t_{-i}) > 0$. Moreover, since, $\mu_i\,(\theta_i, t_{-i}|t_i) = 1$ and $\mu_{-i}\,(\theta_{-i}, t_i|t_{-i}) = 1$, it follows that

$$\frac{\eta(\theta_i, t_i, t_{-i})}{\eta_i(t_i)} = 1 \qquad \text{and} \qquad \frac{\eta(\theta_{-i}, t_i, t_{-i})}{\eta_{-i}(t_{-i})} = 1. \tag{33}$$

The first equality above implies that $\eta((\Theta \times \{t_i\} \times T_{-i}') \setminus \{(\theta_i, t_i, t_{-i})\}) = 0$. In particular,

$$\eta\,(\{(\theta_{-i}, t_i, t_{-i})\}) = 0. \tag{34}$$

The second equality in (33) implies that

$$\eta\left(\left(\Theta \times T_i' \times \{t_{-i}\}\right) \setminus \{(\theta_{-i}, t_i, t_{-i})\}\right) = 0. \tag{35}$$

It follows from (34) and (35) that $\eta(\Theta \times T_i \times \{t_{-i}\}) = 0$. This contradicts our earlier conclusion that $\eta_{-i}(t_{-i}) > 0$. Thus, u_i is not a common prior type.

Being contained in the finitely shy set T_i^0, the set of common prior types is also finitely shy.

References

Anderson, Robert M. and William R. Zame (2001), "Genericity with infinite many parameters." *Advances in Theoretical Economics*, 1, 1–62.

Aumann, Robert J. (1976), "Agreeing to disagree." *Annals of Statistics*, 4, 1236–1239.

Aumann, Robert J. (1987), "Correlated equilibrium as an expression of bayesian rationality." *Econometrica*, 55, 1–18.

Bergemann, Dirk and Stephen Morris (2009), "Robust implementation in direct mechanisms." *Review of Economic Studies*, 76, 1175–1204.

Bernheim, B. Douglas (1984), "Rationalizable strategic behavior." *Econometrica*, 52, 1007–1028.

Billingsley, P. (1999), *Convergence of Probability Measures*. John Wiley & Sons, 2nd edition.

Brandenburger, Adam and Eddie Dekel (1987), "Rationalizability and correlated equilibria." *Econometrica*, 55, 1391–1402.

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

Chen, Yi-Chun, Alfredo Di Tillio, Eduardo Faingold, and Siyang Xiong (2010), "Uniform topologies on types." *Theoretical Economics*, 5, 445–478.

Christensen, Jens P. R. (1974), *Topology and Borel Structure*, volume 10 of *North-Holland Mathematical Studies*. North-Holland, Amsterdam.

Dekel, Eddie, Drew Fudenberg, and Stephen Morris (2006), "Topologies on types." *Theoretical Economics*, 1, 275–309.

Dekel, Eddie, Drew Fudenberg, and Stephen Morris (2007), "Interim correlated rationalizability." *Theoretical Economics*, 2, 15–40.

Ely, Jeffrey C. and Marcin Pęski (2011), "Critical types." *Review of Economic Studies*, forthcoming.

Harsanyi, John (1967-68), "Games with incomplete information played by bayesian players, i-iii." *Management Science*, 14, 159–182, 320–334, 486–502.

Heifetz, Aviad and Zvika Neeman (2006), "On the generic (im)possibility of full surplus extraction in mechanism design." *Econometrica*, 74, 213–234.

Kajii, Atsushi and Stephen Morris (1998), "Payoff continuity in incomplete information games." *Journal of Economic Theory*, 82, 267–276.

Lipman, Barton (2003), "Finite order implications of common priors." *Econometrica*, 71, 1255–1267.

Mertens, Jean-Francois and Shmuel Zamir (1985), "Formulation of bayesian analysis for games with incomplete information." *International Journal of Game Theory*, 14, 1–29.

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

Monderer, Dov and Dov Samet (1996), "Proximity of information in games with incomplete information." *Mathematics of Operations Research*, 21, 707–725.

Neeman, Zvika (1996), "Approximating agreeing to disagree results with common p-beliefs." *Games and Economic Behavior*.

Pearce, David G. (1984), "Rationalizable strategic behavior and the problem of perfection." *Econometrica*, 52, 1029–1050.

Rubinstein, Ariel (1989), "The electronic mail game: A game with almost common knowledge." *American Economic Review*, 79, 389–391.

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