From Anchors to Hierarchies: Identifying Levels of Reasoning in Games*

Adam Brandenburger† Amanda Friedenberg‡ Terri Kneeland§

September 22, 2025

Abstract

We examine the extent to which level-k analysis provides evidence of bounded reasoning in games. Our opening example exhibits a game in which, regardless of the anchor used to initiate the level-k reasoning process, a particular strategy is at most level 1. At the same time, this strategy is consistent with rationality and $(m-1)$th-order belief in rationality (R$(m-1)$BR) for all integers m. The example suggests that the categorization of strategies by the level-k model may overestimate the extent of bounded reasoning on the part of players. However, there is an implicit identification assumption underneath this claim, albeit one standard in epistemic game theory. This is the assumption of a "rich" epistemic type structure that encodes many (sometimes, all) hierarchies of beliefs for the players. It is this assumption that allows us to say that play of an m-rationalizable strategy is consistent with R$(m-1)$BR. Our next result concerns the implications of a suitable identification assumption for level-k analysis. Specifically, we build what we call a "complete level-k type structure" – that encodes the presence of an anchor on which players build their hierarchies of beliefs, and no further restriction. One might conjecture that, in this case, the condition of R$(m-1)$BR would isolate the level-m strategies. Our main theorem shows that this is false: R$(m-1)$BR in a complete level-k type structure once again returns all m-rationalizable strategies. Finally, we find an additional identifying assumption under which epistemic analysis does deliver level-k strategies, and we also assess the verifiability of this assumption.

The Level-k (Nagel, 1995; Stahl and Wilson, 1994, 1995; Costa-Gomes, Crawford and Broseta, 2001; Costa-Gomes and Crawford, 2006) and the related cognitive hierarchy (Camerer, Ho and Chong, 2004) models have played an instrumental role in behavioral game theory. They have gained prominence precisely because of their ability to explain departures from equilibrium in both experimental data and in applications. At the same time, these models have come to serve as a lens through which experimenters have assessed players' reasoning—and bounded reasoning—in games.

This paper revisits the claim that the categorization of levels, as offered by the level-k literature, can provide *direct* information about how players reason—be it reasoning about rationality, reasoning about

*An earlier version was circulated under the title "Two Approaches to Iterated Reasoning in Games" (December 2020). We thank Marciano Siniscalchi for helpful conversations and many seminar audiences for helpful feedback. Brandenburger acknowledges financial support from NYU Stern School of Business, NYU Shanghai, and J.P. Valles. Kneeland acknowledges financial support from ERC Grant SUExp - 801635.

†Stern School of Business, Tandon School of Engineering, NYU Shanghai, New York University, adam.brandenburger@stern.nyu.edu, adambrandenburger.com

‡Department of Economics, University of Michigan, amanda@amandafriedenberg.org, amandafriedenberg.org

§Department of Economics, University of Calgary, terri.kneeland@ucalgary.ca, www.tkneeland.com

irrationality, reasoning about unsophisticated behavior, depths of reasoning or steps in reasoning. It argues that the current interpretation of the level-k model overestimates the extent to which there is evidence of "bounded reasoning" in experimental data.

To make this point, we focus on the basic level-k model. That analysis begins with what is called an anchor, i.e., an exogenous distribution about how the game is played. The anchor is associated with a distribution of so-called level-0 behavior. A level-1 player has a belief that corresponds to the anchor and plays a best response given that belief. The strategies that are a best response to such a belief correspond to level-1 behavior. A level-2 player has a belief that assigns probability 1 to level-1 behavior and plays a best response given such a belief. And so on.

To understand our approach, begin with a known result:

> **Baseline Result:** Fix an anchor. If there is a $k \geq 1$ so that a strategy is classified as level k for that anchor, then the same strategy is k-rationalizable, i.e., survives k rounds of rationalizability.

See, e.g., Costa-Gomes and Crawford (2006, pp. 1739) and Schipper and Zhou (2024, Proposition 1). As a consequence of this result, if a strategy is classified as level k, then there is an $m \geq k$ so that the strategy is m-rationalizable. (Note, the strategy is k-rationalizable and so m-rationalizable for some $m \geq k$.) Standard results in epistemic game theory establish that a strategy is m-rationalizable if and only if it is consistent with *rationality and* $(m-1)^{th}$*-order belief of rationality* ($R(m-1)BR$). See, e.g., Brandenburger and Dekel (1987) and Tan and Werlang (1988). Thus, if a strategy is classified as level k then there is an $m \geq k$ so that the strategy is consistent with R$(m-1)$BR.

The baseline result points to a preliminary approach for relating the categorization from the level-k model to steps of reasoning about rationality:

> If a strategy is classified as level k and there is no $m > k$ so that the strategy is m-rationalizable, then the strategy is consistent with R$(k-1)$BR but is inconsistent with RmBR for all $m > k$. Thus, a level of k captures the maximum level of reasoning about rationality consistent with the data.

However, there are many examples where a strategy is classified as level k, despite the fact that the strategy is consistent with m-rationalizability for $m > k$. This can occur because the strategy is, in fact, also classified as level $m > k$ for the same anchor. (See, e.g., Example 1 in Schipper and Zhou, 2024.) Or, it can occur because the strategy is classified as level $m > k$ for a different anchor. But, importantly, it can also occur even if, for every possible anchor, the strategy is classified as at most level k. Section 1 provides such an example. The example features a strategy that can be classified as level 1, for an appropriate anchor. However, for any anchor, the strategy cannot be classified as level $k \geq 2$, despite the fact that it is consistent with rationality and common belief of rationality.

This last paragraph already suggests that the categorization given by the level-k model may overestimate the extent of bounded reasoning: If a strategy is consistent with RmBR, then it is consistent with $(m+1)$-steps of reasoning about rationality. But it may also be consistent with $(m+1)$-steps of reasoning about rationality and subsequent steps of reasoning about irrationality. And, similarly, if a strategy is consistent with RmBR, then it is also consistent with $(m+1)$ steps of interactive reasoning, e.g., reasoning through sentences of the form "I think, you think, ..."

That said, this conclusion rests on a particular identification assumption. To better understand the assumption, return to the statement that any m-rationalizable strategy is consistent with R$(m-1)$BR.

There is an important background assumption: that players have a "rich" set of hierarchies of beliefs. The implicit identification assumption is that the analyst cannot rule out hierarchies of beliefs. If the analyst knew that the players themselves rule out certain hierarchies of beliefs, then the predictions of R$(m-1)$BR may well be a strict subset of the m-rationalizable strategies. (See Chapter 7 in Battigalli, Friedenberg and Siniscalchi, 2012 for examples.)

This implicit identification assumption is important for the level-k model. In the level-k model, the analyst deliberately chooses an anchor and admits only hierarchies of beliefs that are faithful to the anchor. The choice of the anchor (and so hierarchies) can rest on substantive arguments, such as which behavior is viewed as salient in a particular setting. Or the analyst may hypothesize that hierarchies are faithful to some anchor and attempt to estimate the anchor. In either case, the analyst hypothesizes that players rule out hierarchies of beliefs inconsistent with the anchor.

With this in mind, we focus on a restricted inference problem, one where the analyst has an auxiliary assumption that hierarchies of beliefs are faithful to an anchor. To formalize this inference problem, we follow the approach in the epistemic game theory literature by modeling players beliefs with an epistemic type structure. We focus on a class of such type structures, which we call level-k type structures. These are type structures where players' hierarchies of beliefs are generated by an anchor. An important level-k type structure is, what we call, a complete level-k type structure. This is a level-k type structure that induces a rich set of beliefs that are consistent with the anchor; in a sense, it is a type structure that does not impose substantive restrictions on beliefs that go above and beyond the restrictions that stem from the anchor. (See Sections 3.2 and 8.A.)

Level-k type structures are engineered to mimic the logic of the level-k model. Despite this fact, in a complete level-k type structure, the predictions of R$(m-1)$BR are exactly the m-rationalizable strategies. (See Theorem 6.1.) Note, this is irrespective of the particular anchor that generates the complete level-k type structure. The result has important implications for the restricted inference problem: For a particular anchor, a strategy can be categorized as level k (but not level $m \geq k+1$); yet, there may be an $m \geq k$ so that the same strategy is consistent with RmBR, even when hierarchies of beliefs are required to be consistent with the very same anchor.

Why is there a disconnect between the R$(m-1)$BR predictions in a complete level-k type structure and the categorization from the level-k analysis? The key is that the level-k model only imposes an exogenous restriction on the players' *partial hierarchies of beliefs*. To better understand what this involves, consider a level-2 player, who has a belief that other players have a belief (about play) that corresponds to the anchor. This is distribution on the set of first-order beliefs—i.e., a distribution on what others believe about the play of the game. A second-order belief, however, is a joint distribution about the strategies and first-order beliefs—i.e., a joint distribution about how others play the game and what they believe about the play of the game. The level-k model obtains the full second-order belief endogenously, through the solution concept. In doing so, it imposes an auxiliary requirement that a player cannot rationalize different strategies played by different first-order beliefs. Indeed, in a complete level-k type structure, there will be types that mimic such level-2 players, called 2-types, and those types will not be able to rationalize different strategies played with different first-order beliefs. However, there will be other types—types that are consistent with the partial hierarchies of beliefs induced by the anchor—which can rationalize different strategies played with different first-order beliefs. That is, by explicitly modeling the hierarchies of beliefs consistent with the anchor, we can see that there is a richer set of m^{th}-order beliefs that are consistent with the anchor.

This raises the question: Are there different epistemic assumptions so that the predictions of round k correspond exactly to the categorization of level k? If so, those assumptions would provide a sense in which the categorization of a subject as level k does correspond to k steps of reasoning. Theorem 6.1 provides an answer in the affirmative. On the plus side, the logic behind the result mimics the logic associated with the level-k model, suggesting that our approach (throughout this paper) is tight. On the other hand, as we will discuss, the epistemic analysis points to an arguably new identification assumption: That is, in concluding that a categorization of level k reflects k steps of reasoning about rationality, the analyst is imposing an additional identification assumption, one that goes beyond the requirement that hierarchies are induced by an anchor. (See Identification Assumption 2.) Importantly, that assumption appears difficult to verify in practice.

Literature This is not the first paper to point to difficulties in drawing inferences about how players reason from the level-k categorization. The literature has pointed to at least four difficulties. First, it may be difficult to ascertain the anchor that generates players' beliefs. Toward that end, some papers have suggested looking for a best-fitting anchor (Crawford and Iriberri, 2007; Wright and Leyton-Brown, 2019) or providing auxiliary evidence on the anchor (Costa-Gomes and Crawford, 2006; Brocas et al., 2014). Second, it may be that the players themselves are uncertain about the anchor. (This is captured by Strzalecki's, 2014, cognitive rationalizability and is in the spirit of Section 2.3.2 in Alaoui and Penta, 2016.) Third, there may be that measurement error or other noise in the data, which may make it difficult to infer a categorization of level k, from observed play. (See Stahl and Wilson, 1995, Costa-Gomes and Crawford, 2006, and Cooper et al., 2024.) Fourth, it may be that the inferred levels on reasoning are not portable across games. (See Georganas, Healy and Weber, 2015, Alaoui and Penta, 2016, Alaoui, Janezic and Penta, 2020, and Cooper, Fatas, Morales and Qi, 2024.)

We abstract from these important concerns and looks at an idealized setting. In particular, it focuses on a setting where there is one anchor that generates players' hierarchies of beliefs and that anchor is known to the analyst. So, neither the players nor the analyst face uncertainty about the anchor. Moreover, there is no measurement error or noise in the data. In addition, it ignores concerns about portability. It argues that, even in this idealized setting, the level k categorization may overestimate the extent to which there is bounded reasoning.

The paper sits within a growing literature aimed at bringing ideas from epistemic game theory to bear on experimental data. (Examples include Kneeland, 2015, Ghosh, Heifetz and Verbrugge, 2016, Ghosh and Verbrugge, 2018, Li and Schipper, 2020, Brandenburger, Danieli and Friedenberg, 2021, Friedenberg and Kneeland, 2024, and Healy, 2024.) Moreover, it can be viewed as providing a bridge between the level-k literature and epistemic game theory. Schipper and Zhou (2024) and Liu and Ziegler (2025) are two recent attempts to provide such a bridge. Schipper and Zhou uses ideas from epistemic game theory to motivate a notion of level-k reasoning in extensive-form games. Liu and Ziegler models a level-0 player as one that has different payoffs from those specified in the game; it then uses rationalizability concepts to analyze that game of incomplete information and to draw connections to the level-k literature. Neither paper directly discusses the identification problem that is the focus of this paper.

In the course of our analysis, we introduce the concept of a level-k type structure. This is a particular epistemic type structure that induces hierarchies of beliefs consistent with the anchor. It differs from other rich type structures, meant to model the level-k and cognitive hierarchy concepts, e.g., Kets (2010), Heifetz and Kets (2018), and Strzalecki (2014). The type structures in Kets (2010) and Heifetz and Kets (2018) capture finite-order beliefs about a primitive set of uncertainty, where the players may face uncertainty about the length of others' finite-order beliefs; the type structure in Strzalecki (2014) captures hierarchies of beliefs about numbers (interpreted as levels). Much like Kets and Heifetz and Kets, our framework directly models beliefs about a primitive set of uncertainty. Unlike these papers, we don't include types with finite-order beliefs or beliefs about finite-levels. This stems from the differences in the questions the sets of papers address. Here, we are interested in understanding the extent to which behavior is consistent with high levels of reasoning; as a consequence, being able to rationalize the behavior with a type structure that induces hierarchies of beliefs (as opposed to finite-order beliefs) is a plus. (See, also, the discussion in Section 8.)

1 Heuristic Treatment

Consider the two-player common interest game in Figure 1.1, where Player 1 is denoted by P1 and Player 2 is denoted by P2. We begin by applying the standard level-k solution concept to the game.

		P2			
		a_2	b_2	c_2	d_2
P1	a_1	.9, .9	1, 0	4, 1	1, 0
	b_1	0, 1	4, 4	1, 0	4, 0
	c_1	1, 4	0, 1	0, 0	0, 3
	d_1	0, 1	0, 4	3, 0	3, 3

Figure 1.1: A Common-Interest Game

The level-k solution concept begins by fixing an exogenous anchor for each player. For Pi=P1,P2 this is a distribution μ_i on the strategies the other player, Pj, can choose. The level-1 strategies for Pi are the strategies that are a best response under μ_i. The level-2 strategies for Pi are the strategies that are a best response under a belief that assigns probability 1 to level-1 strategies of Pj. And so on.

Figure 1.2 describes the level-k behavior in four examples. In each example, P1 and P2 have the same anchor, i.e., $\mu_1 = \mu_2$: This is either the uniform anchor, the anchor where Pi assigns probability 1 to Pj choosing a_j, the anchor where Pi assigns probability 1 to Pj choosing c_j, or the anchor where Pi assigns probability 1 to Pj choosing d_j. Notice, for each strategy $s_i \in \{a_i, b_i, c_i\}$ and each number $m \geq 1$, there is some anchor so that so that s_i is level-m for Pi.[1]

In each of these examples, there is no m so that d_i is level-m for Pi. If Pi has an anchor that assigns $.5 : .5$ to $c_j : d_j$, then d_i would be level-1. But, regardless of Pj's anchor, d_i cannot be be level-2. More generally:

Claim 1.1. *Suppose P1's and P2's anchors are given by* (μ_1, μ_2)*. If* d_i *is level-*m *for P*i*, then* $m = 1$.

The key observation is that d_i is optimal only under a distribution that assigns positive probability to both c_j and d_j.[2] Therefore, if d_i is level-2 for Pi, it must be that both c_j and d_j are level-1 for Pj. However,

[1] As standard, we refer to the solution concept as "level-k." We use the index m to refer to a particular realization of k.

[2] When Pj is restricted to play a strategy in $\{a_j, b_j, d_j\}$ (resp. $\{a_j, b_j, c_j\}$), d_i is dominated by a mixture of $a_i : b_i$ (resp. by a_i).

	Uniform	**1 to a_j**	**1 to c_j**	**1 to d_j**
Level-1	b_i	c_i	a_i	b_i
Level-2	b_i	a_i	c_i	b_i
Level-3	b_i	c_i	a_i	b_i
Level-4	b_i	a_i	c_i	b_i
Level-5	b_i	c_i	a_i	b_i
...	...	...	...	...

Figure 1.2: Level-k

there is no anchor μ_j under which c_j and d_j are both a best response.[3] Thus, d_i cannot be level-2 for Pi, regardless of Pj's anchor μ_j. And, similarly, d_j is not level-2 for Pj, regardless of Pi's anchor μ_i. This, in turn, implies that d_i is not level-3 for Pi. And so on, for any $m \geq 3$.

The Basic Inference Problem To recap: The strategies d_1 and d_2 are level-1 for some anchor. But, for any anchor and any $m \geq 2$, d_1 and d_2 are not level-m.

Suppose the analyst only observes data about how the game is played (and not auxiliary data, say, about players beliefs). In particular, suppose the analyst observes P1 play d_1. What can the analyst infer about how she reasons? Based on the level-k analysis, the analyst might be tempted to conclude that P1 is rational—in the sense that she plays a best response to the anchor—but does not reason further. Language used in the literature is that P1 believes believes P2 is nonstrategic, P1 reasons one step, or P1 has depth of reasoning one.

However, in this game, the entire strategy set is rationalizable. Standard results in epistemic game theory show that any rationalizable strategy is consistent with rationality and common belief of rationality. (See, e.g., Brandenburger and Dekel, 1987 and Tan and Werlang, 1988.) Thus, the observation of d_1 does not, in and of itself, indicate that a P1 must believe P2 is not strategic.

More generally, the observation of d_1 alone cannot point to a bound in the steps or depth of reasoning, i.e., how many steps of "I think, you think, I think ..." P1 can perform: P1 can only engage in rationality and $(m-1)$ rounds of reasoning about rationality, if she can engage if m-steps of "I think, you think, I think ..." Thus, if behavior is consistent with rationality and common belief of rationality then it is also consistent with an unbounded depth of reasoning.

Rationality and Common Belief of Rationality It will be useful to better understand what goes into the statement that d_1 is consistent with rationality and common belief of rationality. To better understand, we revisit a standard epistemic model, as applied to Figure 1.1. A hallmark of the model is that it describes the players hierarchies of beliefs about the play of the game. This is a necessary step: To specify whether P1 (resp. P2) is rational, we must describe what beliefs P1 (resp. P2) holds about P2's (resp. P1's) play. After all, whether a strategy is a best response for P1 depends on these first-order belief. By a similar logic, to specify whether P1 does or does not believe P2 is rational, we must describe P1's joint belief about

[3] If c_j and d_j have the same expected payoff, then the expected payoff of a_j must be strictly higher.

P2's strategy and first-order belief, i.e., about P2's strategy and belief about P2's play. After all, whether a strategy is rational or irrational for P2 will depend on his first-order belief. And so on.

We model these hierarchies of beliefs by an epistemic type structure, in the spirit of Harsanyi (1967). The type structure has two ingredients: First, for each Pi, there is a set of types T_i; in our example,

$$T_i = \{t_i, u_i, v_i, w_i\}.$$

Second, for each Pi, there is a belief map β_i, which maps each type of Pi to a belief about the strategy-type pairs of Pj; in our example

$$\beta_i(t_i)(c_j, v_j) = 1 \qquad \beta_i(u_i)(b_j, u_j) = 1 \qquad \beta_i(v_i)(a_j, t_j) = 1$$

and

$$\beta_i(w_i)(c_j, v_j) = \beta_i(w_i)(d_j, w_j) = \frac{1}{2}.$$

Each type induces hierarchies of beliefs about the play of the game. For instance, type t_i assigns probability 1 to Pj playing c_j, while type v_i assigns assigns probability 1 to Pj playing a_j. Since t_i assigns probability 1 to (c_j, v_j), this implies that t_i assigns probability 1 to "Pj plays c_j and believes I play a_i." And so on. See Section 2 for more details.

Now turn to rationality, belief in rationality, etc. Rationality is a property of a strategy-type pair. The pair (a_i, t_i) is rational because a_i maximizes Pi's expected payoffs given the belief associated with t_i: The action a_i is a best response to c_j. In fact, the set of rational strategy-type pairs for Pi is

$$R_i = \{(a_i, t_i), (b_i, u_i), (c_i, v_i), (d_i, w_i)\}.$$

Now observe that each type of Pi assigns probability 1 to "Pj is rational," i.e., to the event R_j; thus, each type of Pi believes the other player is rational. So, R_i is also the set of strategy-type pairs for Pi that are consistent with rationality and 1^{st}-order belief of rationality. From here, we can iterate to conclude that R_i is, in fact, the set of strategy-type pairs consistent with rationality and common belief of rationality (RCBR). As a consequence, each of a_i, b_i, c_i, and d_i are consistent with RCBR.

Hierarchies of Beliefs vs. Anchored Beliefs We have seen that the strategy d_i is, in fact, consistent with RCBR. To show this, we described a specific model of P1's and P2's hierarchies of beliefs and pointed to a type in that model, namely w_i, so that (d_i, w_i) is rational, believes Pj is rational, and so on. Importantly, those hierarchies of beliefs were inconsistent with the idea that the players' hierarchies are generated by an anchor. Take, for instance, the case where P1's and P2's anchors (μ_1, μ_2) both assign probability one to the other player Pj choosing c_j. Type v_1 has the first-order belief associated with P1's anchor μ_1 and type t_1 believes P2 has the first-order belief associated with P2's anchor μ_2. But, types u_1 and w_1 do not have hierarchies of beliefs consistent with these anchor. Similarly, if P1's and P2's anchors (μ_1, μ_2) both assign .5 : .5 to the other player Pj choosing c_j or d_j. Then, type w_1 has first-order beliefs associated with P1's anchor. But no other type has hierarchies consistent with this anchor. And so on. (See Example 4.1.)

Arguably, the spirit of level-k analysis involves a restriction on the hierarchies of beliefs that players can hold. In particular, the analysis imposes the substantive assumption that the players beliefs are generated by a particular anchor. This assumption is important in categorizing a particular strategy as level-m for some $m \geq 1$.

This raises the question: Suppose players hierarchies of beliefs are generated by an anchor. In that case, would the observation of d_i allow us the analyst to conclude that Pi's behavior is inconsistent with Pi being rational and believing Pj is rational? That is, would the observation of d_i point to a form of bounded reasoning?

The Restricted Inference Problem To address the question, our analysis focuses on, what we call, (epistemic) level-k type structures. Much as above, these are type structures that involve type sets and belief maps for each of P1 and P2. But, now, the type set of Pi can be decomposed into a set of 1-types (T_i^1), a set of 2-types (T_i^2), etc. The 1-types each have first-order beliefs associated with the anchor. The 2-types each assign probability one to Pj having a 1-type (i.e., their marginal belief on T_j assigns probability 1 to T_j^1). And so on. Notice, a level-k type structure is defined relative to a particular anchor and only induces hierarchies of beliefs consistent with that anchor. (See Proposition 4.1.) Thus, the type structure cannot induce all hierarchies of beliefs.

A notable level-k type structure is, what we call, a *complete level-k type structure*. This is a level-k type structure that satisfies the following requirement: For every belief that assigns probability 1 to the m-types of Pj, there is an $(m+1)$-type of Pi that holds induces that belief. A complete level-k type structure induces a rich set of beliefs. (See Proposition 4.3 and Section 8-A.) Proposition 4.3 shows that there exists a complete level-k type structure.

The main theorem provides the behavioral implications of rationality and m^{th}-order belief of rationality (RmBR) in a complete level-k type structure.

> **Main Theorem (Theorem 5.1).** In a complete level-k type structure (for a particular anchor), the predictions of RmBR are exactly the $(m+1)$-rationalizable strategies.

Thus, even when we focus on models of hierarchies of beliefs that are consistent with the anchor, each $(m+1)$-rationalizable strategy is consistent with RmBR.

Return then to Figure 1.1. If we observe P1 play d_1 we cannot conclude that there is a bound m so that the behavior must reflect RmBR, even if we assume that the hierarchies of beliefs are generated by a particular anchor. Thus, the categorization of d_1 as level-1 does not allow us to draw a conclusion about bounded reasoning—at least not without additional auxiliary assumptions about how players reason or without a richer dataset. Section 6 discusses additional auxiliary assumptions and the difficulty of verifying those assumptions in the data.

2 The Environment

We begin with mathematical preliminaries used throughout the paper. Fix a metrizable set Ω and endow Ω with the Borel σ-algebra. We will refer to an element of the Borel σ-algebra as an *event*. Write $\Delta(\Omega)$ for the set of Borel probability measures on Ω and endow $\Delta(\Omega)$ with the topology of weak convergence. Given a measure $\mu \in \Delta(\Omega \times \Phi)$, write $\text{marg}_{\Omega}\mu$ for the marginal of μ on Ω.

Given a finite index set I and a collection of metrizable sets $(\Omega_i : i \in I)$, write $\Omega_{-i} = \prod_{j \in I \setminus \{i\}} \Omega_j$ and $\Omega = \prod_{i \in I} \Omega_j$. Endow the product of metrizable spaces with the product topology. Given a second collection of metrizable sets $(\Phi_i : i \in I)$ and measurable maps $f_i : \Omega_i \to \Phi_i$, write $f_{-i} = \Omega_{-i} \to \Phi_{-i}$ for the associated product map, i.e., given $\omega_{-i} = (\omega_j : j \neq i)$, $f_{-i}(\omega_{-i}) = (f_j(\omega_j) : j \neq i)$. If each f_i is measurable (resp. continuous), then each f_{-i} is also measurable (resp. continuous).

Fix metrizable sets Ω and Φ and let $f : \Omega \to \Phi$ be a measurable map. The image measure of f under $\mu \in \Delta(\Omega)$ is a measure $\nu \in \Delta(\Phi)$ where, for each Borel $E \subseteq \Phi$, $\nu(E) = \mu(f^{-1}(E))$. Let $\underline{f} : \Delta(\Omega) \to \Delta(\Phi)$ map each $\nu \in \Delta(\Omega)$ to the image measure of f under ν. Note, $\underline{f}$ is measurable; if f is continuous, $\underline{f}$ is continuous. (See, e.g., Friedenberg and Keisler, 2021, Lemma A.1, and Aliprantis and Border, 2007, Theorem 14.14.)

2.1 The Epistemic Game

Throughout the paper, fix a game $G = (S_i, \pi_i : i \in I)$: Here, I is a finite set of players, S_i is a finite strategy set for player i, and $\pi_i : S_i \times S_{-i} \to \mathbb{R}$ is player i's payoff function. The game is non-trivial, in that each player has at least two strategies ($|S_i| \geq 2$). Extend π_i to $\pi_i : S_i \times \Delta(S_{-i}) \to \mathbb{R}$ in the usual way.

An epistemic game appends to the game a description of the players' hierarchies of beliefs about the play of the game. Following Harsanyi (1967), we use type structures to implicitly describe the hierarchies of beliefs.

Definition 2.1. An S**-based type structure** is some $\mathcal{T} = (S_{-i}, T_i, \beta_i : i \in I)$ where,

(i) for each i, T_i is a metrizable set of **types** for i, and

(ii) for each i, $\beta_i : T_i \to \Delta(S_{-i} \times T_{-i})$ is a measurable **belief map** for i.

In an S-based type structure, each type of player i, t_i, is mapped to a joint belief about the strategies and types of the other players. Because the set of strategies is fixed throughout our analysis, we often refer to an S-based type structure as, simply, a **type structure**. When each T_i is (at most) countable, we call the type structure **countable**.

2.2 Type Structures and Hierarchies of Beliefs

The epistemic game describes the rules of the game, payoff functions, and hierarchies of beliefs about the play of the game. The former two are captured by G and the latter is captured by a type structure. The example in Section 1 is indicative of how types induce hierarchies of beliefs. In particular, each type $\tilde{t}_i$ induces a belief on the strategies of other players, given by $\text{marg}_{S_{-i}} \beta_i(\tilde{t}_i)$. For instance, type w_i's first-order belief assigns $.5 : .5$ to $c_j : d_j$. Moreover, because each type has a joint belief about the strategies and types of the other player, each type has a joint belief about the strategies and first-order beliefs of the other player. For instance, type w_i assigns .5 to "the other player will play c_j and believes that I will play a_i" and .5 to "the other player will play d_j and assigns $.5 : .5$ to me playing $c_i : d_i$." These joint beliefs constitute the type's second-order beliefs. And so on.

The example is indicative of how type structures induce hierarchies of beliefs. To formalize this, begin by inductively describing the set of m^{th}-order beliefs of player i. Set $X_i^1 = S_{-i}$ and $H_i^1 = \Delta(X_i^1)$; these sets are each compact metric sets. Assume the sets X_i^m and H_i^m have been defined and are compact metric sets. Set

$$X_i^{m+1} = \{(s_{-i}, h_{-i}^1, \ldots, h_{-i}^m) \in X_i^m \times H_{-i}^m : \text{ if } m \geq 2 \text{ then, for each } j \neq i, \text{ marg}_{X_j^m} h_j^m = h_j^{m-1}\}$$

and $H_i^{m+1} = \Delta(X_i^{m+1})$. These too are compact metric sets. (See Friedenberg, 2010, Lemma A1 and Remark A1.) The set X_i^m is player i's $\boldsymbol{m^{th}}$**-order space of uncertainty**. The set H_i^m is player i's set of $\boldsymbol{m^{th}}$**-order beliefs**. Then

$$H_i^\infty = \{(h_i^1, h_i^2, \ldots) \in \prod_{m\geq 1} H_i^m : \text{ for each } m, \text{marg}_{X_i^m} h_i^{m+1} = h_i^m\}$$

is player i's set of **hierarchies of beliefs**.

For each $m \geq 1$, there is a natural mapping $\delta_i^m : T_i \to H_i^m$, specifying each type's m^{th}-order belief. Type t_i's first-order belief is simply the marginal of $\beta_i(t_i)$ onto the strategies of the other players; that is, $\delta_i^1(t_i) = \text{marg}_{S_{-i}} \beta_i(t_i)$. Type t_i's second-order belief, $\delta_i^2(t_i) = h_i^2$, is a joint belief about strategies and first-order beliefs: The probability that h_i^2 assigns to an event about $S_{-i} \times H_{-i}^1$ corresponds to the probability that $\beta_i(t_i)$ assigns to strategy-type pairs that induce that event. More precisely, for each event $E_{-i} \subseteq X_i^1 \times H_{-i}^1 = S_{-i} \times \prod_{j\neq i} \Delta(S_{-j})$,

$$h_i^2(E_{-i}) = \beta_i(t_i)(\{(s_{-i}, t_{-i}) : (s_{-i}, \delta_{-i}^1(t_{-i})) \in E_{-i}\}).$$

Appendix B.1 formally describes the maps $\delta_i^m : T_i \to H_i^m$. Given these maps, $\delta_i : T_i \to H_i^\infty$ is defined by $\delta_i(t_i) = (\delta_i^1, \delta_i^2, \ldots)$. If $\delta_i(t_i) = h_i$ (resp. $\delta_i^m(t_i) = h_i^m$), say that type t_i **induces** the hierarchy of beliefs h_i (resp. the m^{th}-order belief h_i^m). The set of hierarchies of beliefs for i **induced by** $\mathcal{T}$ is $\delta_i(T_i) \subseteq H_i$.

Of particular interest is a type structure that is "rich," in the sense that it induces all possible beliefs.

Definition 2.2. Call the type structure $\mathcal{T} = (S_{-i}, T_i, \beta_i : i \in I)$ **type-complete** if, for each i, β_i is onto.

So, a type structure is type-complete if, for each belief i can hold (about $S_{-i} \times T_{-i}$), there is a type of i that holds that belief. That is, the type structure contains all possible beliefs about types. The canonical constructions of a so-called universal type structure (e.g,. Mertens and Zamir, 1985, Brandenburger and Dekel, 1993, Heifetz and Samet, 1998, etc) are each type-complete. When the type sets are compact and the belief maps are continuous, a type-complete type structure induces all hierarchies of beliefs. (See Friedenberg, 2010.)

3 Hierarchies of Beliefs Induced by the Anchor

The level-k solution concept is tied to an **anchor** $\boldsymbol{\mu} = (\mu_i : i \in I) \in \prod_{i\in I} \Delta(S_{-i})$; call μ_i $\boldsymbol{i}$**'s anchor**. Conceptually, an anchor specifies a first-order belief for each player i. This implicitly limits the hierarchies of beliefs the players consider possible. However, importantly, the anchor alone does not uniquely pin down those hierarchies. Instead, it restricts, what we will call, the hierarchies of partial beliefs. We next describe how the anchor restricts the partial hierarchies and, in turn, restricts the hierarchies of beliefs.

Remark 3.1. The literature will often fix a symmetric game and look at symmetric anchors, i.e., anchors where each player has the same belief about how others play the game. (There are important exception.) Because we apply the ideas to arbitrary games (i.e., not necessarily symmetric games), we do not restrict the anchors to be symmetric. To be sure, players' anchors can be symmetric, but they need not be symmetric. Likewise, anchors can involve a belief that is independent or correlated. They can involve degenerate or non-degenerate beliefs. Etc.

3.1 Hierarchies of Partial Beliefs

Set $P_i^1 = \Delta(S_{-i})$ and note that it is a compact metric set. Assuming compact metric sets P_i^m have been defined, set $P_i^{m+1} = \Delta(P_{-i}^m)$ and note that it too is a compact metric set. The set P_i^m is player i's set of $\boldsymbol{m^{th}}$**-order partial beliefs**. Note, when $m \geq 2$, an m^{th}-order partial belief differs from an m^{th}-order belief. For instance, a second-order belief is a joint belief about strategies and first-order beliefs, whereas a second-order partial belief is a belief only about first-order beliefs. Write

$$P_i^\infty = \prod_{m \geq 1} P_i^m$$

for the set of **hierarchies of partial beliefs**.

The anchor implicitly imposes a restriction on the m^{th}-order partial beliefs that players consider possible. For instance, if i is a level-1 player, then i's first-order partial belief must correspond to the anchor. If i is a level-2 player, then i's second-order partial belief must assign probability 1 to the first-order beliefs $\boldsymbol{\mu}_{-i} := (\mu_j : j \in I \backslash \{i\})$. And so on.

More generally, an anchor $\boldsymbol{\mu} = (\mu_i : i \in I) \in \prod_{i \in I} \Delta(S_{-i})$ uniquely determines m^{th}-order partial beliefs, $p_{i,\boldsymbol{\mu}}^m$: Set $p_{i,\boldsymbol{\mu}}^1 = \mu_i$. Assuming each $p_{i,\boldsymbol{\mu}}^m \in P_i^m$ has been defined, let $p_{i,\boldsymbol{\mu}}^{m+1} \in P_i^{m+1}$ be the measure with $p_{i,\boldsymbol{\mu}}^{m+1}(\{p_{-i,\boldsymbol{\mu}}^m\}) = 1$. Write $\boldsymbol{p}_{i,\boldsymbol{\mu}} = (p_{i,\boldsymbol{\mu}}^1, p_{i,\boldsymbol{\mu}}^2, \ldots)$ and $\boldsymbol{p}_{\boldsymbol{\mu}} = (p_{\boldsymbol{\mu}}^1, p_{\boldsymbol{\mu}}^2, \ldots)$.

3.2 Hierarchies of Beliefs Consistent with the Anchor

There is a natural mapping from hierarchies of beliefs to hierarchies of partial beliefs, viz. $\eta_i : H_i^\infty \to P_i^\infty$. To understand the mapping, consider $\eta_i(h_i^1, h_i^2, \ldots) = (p_i^1, p_i^2, \ldots)$. Intuitively, $p_i^1 = h_i^1$ since there is no distinction between first-order beliefs and first-order partial beliefs. Moreover, $p_i^2 = \text{marg}_{\,\Delta(S_{-i})} h_i^2$, since a second-order partial belief simply provides information about beliefs over first-order partial beliefs and first-order partial beliefs correspond to first-order beliefs. Since there is a distinction between second-order partial beliefs and second-order beliefs, the relationship between h_i^3 and p_i^3 requires care.

To define the mapping η_i, it will be convenient to define sets that correspond to i's m^{th}-order space of partial uncertainty, i.e., Y_i^m: Set $Y_i^1 = S_{-i}$ and, for $m \geq 2$, $Y_i^m = P_{-i}^{m-1}$. Note that $P_i^m = \Delta(Y_i^m)$. Now, inductively define continuous maps $\hat{\eta}_i^m : X_i^m \to Y_i^m$ and $\eta_i^m : H_i^m \to P_i^m$: First, take $\hat{\eta}_i^1 : X_i^1 \to Y_i^1$ and $\eta_i^1 : H_i^1 \to P_i^1$ to be the identity maps; note that these are continuous. Next, assume continuous maps $\hat{\eta}_i^m : X_i^m \to Y_i^m$ and $\eta_i^m : H_i^m \to P_i^m$ have been defined. Define $\hat{\eta}_i^{m+1} : X_i^{m+1} \to Y_i^{m+1}$ so that, for each $x_i^{m+1} = (x_i^m, h_{-i}^m) \in X_i^{m+1}$, $\hat{\eta}_i^{m+1}(x_i^m, h_{-i}^m) = \eta_{-i}^m(h_{-i}^m)$. Since each η_j^m is continuous, $\hat{\eta}_{-i}^{m+1}$ is continuous. Now let $\eta_i^{m+1} = \underline{\hat{\eta}}_i^{m+1}$, so that $\eta_i^{m+1}(h_i^{m+1})$ is the image measure of h_i^{m+1} under $\hat{\eta}_i^{m+1}$; note that $\eta_i^{m+1} = \underline{\hat{\eta}}_i^{m+1}$ is continuous since $\hat{\eta}_i^{m+1}$ is continuous.

The map $\eta_i : H_i^\infty \to P_i^\infty$ is given by $\eta_i(h_i^1, h_i^2, \ldots) = (\eta_i^1(h_i^1), \eta_i^2(h_i^2), \ldots)$. Thus it maps each hierarchy of beliefs to its associated hierarchy of partial beliefs.

Definition 3.1. Say a hierarchy $h_i = (h_i^1, h_i^2, \ldots)$ is **consistent with the anchor** $\boldsymbol{\mu} = (\mu_i : i \in I) \in \prod_{i \in I} \Delta(S_{-i})$ if there exists some $m \geq 1$ so that $\eta_i^m(h_i^m) = p_{i,\boldsymbol{\mu}}^m$.

If $h_i = (h_i^1, h_i^2, \ldots)$ is consistent with the anchor, there is some m^{th}-order belief that coincides with the m^{th}-order partial beliefs induced by the anchor. This captures the restriction on beliefs implicitly imposed by the level-k solution concept. (Note, there, a player classified as level-m has m^{th}-order partial beliefs induced by the anchor, but may not have n^{th}-order partial beliefs induced by the anchor for some $n \neq m$.)

4 Level-k Type Structures

We will be interested in type structures that only induce hierarchies of beliefs consistent with the anchor. This will be captured by a level-k type structure. This section defines such structures.

4.1 Level-k Type Structure

Fix a type structure $\mathcal{T} = (S_{-i}, T_i, \beta_i : i \in I)$. Say $\mathcal{C}_i = \{T_i^m : m = 1, 2, \ldots\}$ is a **Borel cover** of T_i if (i) each T_i^m is a non-empty Borel subset of T_i, and (ii) $\bigcup_{m \geq 1} T_i^m = T_i$. Note, a countably infinite partition of T_i is a Borel cover, if each of its members is Borel. But, a Borel cover need not be a partition.

Definition 4.1. Call $\mathcal{T} = (S_{-i}, T_i, \beta_i : i \in I)$ a **level-k type structure** (**for** $\boldsymbol{\mu} = (\mu_i : i \in I)$) if, for each i, there exists a Borel cover $\mathcal{C}_i = \{T_i^m : m = 1, 2, \ldots\}$ of T_i so that the following hold:

(i) If $t_i \in T_i^1$, then $\text{marg}_{S_{-i}} \beta_i(t_i) = \mu_i$, and

(ii) For each $m \geq 1$, if $t_i \in T_i^{m+1}$, then $\beta_i(t_i)(S_{-i} \times T_{-i}^m) = 1$.

In a level-k type structure, we can decompose each player's types into non-empty sets $T_i^1, T_i^2, \ldots$. We will refer to types in T_i^m as i's ***m*-types**. The 1-types have first-order beliefs associated with the anchor $\boldsymbol{\mu}$. The 2-types assign probability 1 to the 1-types having the first-order beliefs associated with $\boldsymbol{\mu}$. More generally, the $(k+1)$-types assign probability 1 to the m-types.

Example 4.1. To better understand what goes into a level-k type structure, return to the example of Section 1 (page 7). That type structure is not a level-k type structure for any anchor $\boldsymbol{\mu} = (\mu_1, \mu_2)$. Suppose, contra hypothesis, that this type structure is a level-k type structure for some anchor $\boldsymbol{\mu}$. Then, for each i, there exists some m so that u_i is an m-type. This implies that there must be some player i for which $u_i \in T_i^1$ and, so, $\mu_i(b_j) = 1$. As a consequence, u_i is the unique 1-type for i. If $T_1^1 = \{u_1\}$ and $T_2^1 = \{u_2\}$ then, for each i and each m, $T_i^m = \{u_i\}$. That is, types t_i, v_i, w_i are not m-types for any m. So, without loss of generality, suppose $T_1^1 = \{u_1\}$ and $u_2 \notin T_2^1$. Since $T_1^1 = \{u_1\}$ it follows that

$$T_1^{2m+1} = \{u_1\} \qquad \text{and } T_2^{2(m+1)} = \{u_2\}$$

for each $m \geq 0$.

Observe, since t_2, v_2, and w_2 have distinct first-order beliefs, T_2^1 must be a singleton. Since each $\tilde{t}_1 \in T_1^2$ must assign probability one to T_2^1, T_1^2 must also be a singleton. Now, by induction, for each i and each m, T_i^m must be a singleton. But then, for each i and each m, w_i is not an m-type.

Thus, there can be no anchor $\boldsymbol{\mu}$ so that the example is classified as a level-k type structure for $\boldsymbol{\mu}$. Indeed, an analogous argument shows that there is no $\boldsymbol{\mu}$ so that the type structure only induces hierarchies of beliefs consistent with $\boldsymbol{\mu}$.

This argument reflects the fact that, in the example, the type structure induces hierarchies of beliefs that are inconsistent with a single anchor. By contrast, level-k type structures only induce hierarchies of beliefs consistent with an anchor.

Proposition 4.1. *Let $\mathcal{T}$ be a level-k type structure for $\boldsymbol{\mu}$. Then each hierarchy of beliefs induced by $\mathcal{T}$ is consistent with $\boldsymbol{\mu}$.*

Appendix B.2 proves Proposition 4.1. The proof follows from a stronger claim: If a type is classified as an m-type (according to any appropriately chosen cover), then the type must induce the m^{th}-order partial beliefs $p^m_{i,\boldsymbol{\mu}}$. This provides an interpretation of the m-types.

Because there are (always) hierarchies of beliefs that are inconsistent with the anchor, Proposition 4.1 implies that no level-k type structure induces all hierarchies of beliefs. It is also the case that a level-k type structure cannot induce all beliefs about types.[4]

Proposition 4.2. *If $\mathcal{T} = (S_{-i}, T_i, \beta_i : i \in I)$ is a level-k type structure for $\boldsymbol{\mu}$, then $\mathcal{T}$ is not type-complete.*

To understand why the result holds, suppose that a level-k type structure were, in fact, type-complete. The key is that we can inductively find types t^m_i so that t^m_i is an m-type but not an ℓ-type for any $\ell \neq m$. Now consider a belief $\nu_j \in \Delta(S_{-j} \times T_{-j})$ which (i) has $\text{marg}_{S_{-j}} \nu_j \neq \mu_j$, and (ii) assigns positive probability to both t^1_i and t^2_i. On the one hand, type-completness requires that there is a type of player j that holds that belief. On the other hand, that type cannot be classified as an m-type for any m: Because first-order beliefs differ from the anchor, it cannot be classified as a 1-type. Because it assigns positive probability to both t^1_i and t^2_i but there is no ℓ with $\{t^1_i, t^2_i\} \subseteq T^\ell_i$, it cannot be classified as a m-type for any $m \geq 2$. This results in a contradiction.

4.2 Hierarchies Induced by Level-k Type Structures

While a level-k type structure must induce hierarchies of beliefs consistent with the anchor $\boldsymbol{\mu}$, two different level-k type structures (for $\boldsymbol{\mu}$) may induce different hierarchies of beliefs. The next two examples illustrate this fact.

Example 4.2. Consider a two-player game where each $S_i = \{\Box_i, \Diamond_i\}$. Suppose the anchor $\boldsymbol{\mu} = (\mu_1, \mu_2)$ is such that, for each i, $\mu_i(\Box_{-i}) = \frac{2}{3}$. Consider a type structure $\mathcal{T}$ with the following properties: Set $T_1 = T_2 = \mathbb{N}^+$. Take each $\beta_i(1)$ so that $\beta_i(1)(\Box_{-i}, 2) = \frac{2}{3}$ and $\beta_i(1)(\Diamond_{-i}, 3) = \frac{1}{3}$. For $m \geq 2$, take

$$\beta_i(m)(\Box_{-i}, m-1) = 1 \qquad \text{if } m \text{ is even,}$$

and

$$\beta_i(m)(\Diamond_{-i}, m-1) = 1 \qquad \text{if } m \text{ is odd.}$$

For each i, $\{T^m_i = \{m\} : m \geq 1\}$ is a Borel cover of T_i. Thus, $\mathcal{T}$ is a level-k type structure.

Example 4.3. Consider a two-player game where each $S_i = \{\Box_i, \Diamond_i\}$. Suppose the anchor $\boldsymbol{\mu} = (\mu_1, \mu_2)$ is such that, for each i, $\mu_i(\Box_{-i}) = \frac{2}{3}$. Consider a type structure $\mathcal{T}$ with the following properties: Set $T_1 = T_2 = \mathbb{N}^+$. Likewise, take each $\beta_i(1)$ so that $\beta_i(1)(\Box_{-i}, 2) = \frac{2}{3}$ and $\beta_i(1)(\Diamond_{-i}, 3) = \frac{1}{3}$. For $m \geq 2$, take

$$\beta_i(m)(\Box_{-i}, m-1) = \beta_i(m)(\Diamond_{-i}, m-1)) = \frac{1}{2}.$$

For each i, $\{T^m_i = \{m\} : m \geq 1\}$ is a Borel cover for $\mathcal{T}$. Thus, $\mathcal{T}$ is a level-k type structure.

[4] Friedenberg (2010) shows that a type-complete structure induces all hierarchies of beliefs, if the type sets are compact and the belief maps are continuous. So, there can be no level-k type structure that is type-complete, has compact type sets and continuous belief maps. This result does not impose the results of compactness and continuity, which (as argued in (Friedenberg, 2010)) are really substantive assumptions.

Examples 4.2-4.3 provide two different level-k type structures for a given anchor $\boldsymbol{\mu}$. In both type structures, the 1-types have first-order (partial) beliefs associated with the anchor, i.e., they assign $\frac{2}{3} : \frac{1}{3}$ to $\square_{-i} : \diamond_{-i}$. Likewise, in both type structures, the 2-type have second-order partial beliefs associated with the anchor, i.e., the type $t_i = 2$ assigns probability 1 to $t_{-i} = 1$ and so probability 1 to the event that "the other player assigns $\frac{2}{3} : \frac{1}{3}$ to $\square_{-i} : \diamond_{-i}$." And so on. In this sense, the types induce hierarchies of partial beliefs consistent with the anchor, illustrating Proposition 4.1.

However, in these two examples, the type structures induce disjoint sets of hierarchies of beliefs. To see this, observe that the first-order beliefs of m-types differs in these type structures, when $m \geq 2$. In Example 4.2, each such m-type has a degenerate belief, assigning probability 1 to either of $\square_{-i}$ or $\diamond_{-i}$; in Example 4.3, each such m-type has a non-degenerate belief, assigning $\frac{1}{2} : \frac{1}{2}$ to $s_{-i} : r_{-i}$. Thus, for each type $m \geq 2$ in Example 4.2, there is no type $n \geq 1$ in Example 4.3 that induces the same first-order beliefs, a fortiori the same hierarchies of beliefs. And conversely, with Example 4.3 and Example 4.2 reversed. Moreover, the 1-types induce distinct second-order beliefs. In Example 4.2, type 1 assigns probability $\frac{2}{3}$ to "the other player chooses $\square_{-i}$ and assigns probability 1 to me choosing $\square_i$;" however, in 4.3, type 1 assigns zero probability to that same event.

4.3 Complete Level-k Type Structures

Proposition 4.2 says that a level-k type structure imposes the substantive assumption that the hierarchies are induced by the anchor. But, Section 4.2 illustrated that there may be multiple level-k type structures, associated with the same anchor, but which induce different hierarchies of beliefs. To understand why this arises, note that, in Examples 4.2-4.3 there is exactly one 2-type. Yet, there are many second-order beliefs that a player can hold, even if the player has a second-order partial belief consistent with the anchor. Both type structures rule out such second-order beliefs and, in doing so, they impose auxiliary assumptions on players' hierarchies of beliefs. These auxiliary assumptions on beliefs go above and the substantive assumptions imposed by the anchor. We will be interested in type structures that don't impose these exogenous restrictions on beliefs (or, at least, minimize such exogenous restrictions).

Definition 4.2. Call $\mathcal{T} = (S_{-i}, T_i, \beta_i : i \in I)$ a **complete level-k type structure (for $\boldsymbol{\mu} = (\mu_i : i \in I)$)** if, for each i, there exists a Borel cover $\mathcal{C}_i = \{T_i^m : m = 1, 2, \ldots\}$ of T_i so that the following hold:

(i) If $t_i \in T_i^1$, then $\text{marg}_{S_{-i}} \beta_i(t_i) = \mu_i$,

(ii) For each $m \geq 1$, if $t_i \in T_i^{m+1}$, then $\beta_i(t_i)(S_{-i} \times T_{-i}^m) = 1$, and

(iii) For each $m \geq 1$ and each $\nu_i \in \Delta(S_{-i} \times T_{-i})$ with $\nu_i(S_{-i} \times T_{-i}^m) = 1$, there is a type $t_i \in T_i^{m+1}$.

Call $\mathcal{T}$ a **complete level-k type structure** if there is some $\boldsymbol{\mu}$ so that $\mathcal{T}$ is a complete level-k type structure for μ.

So, $\mathcal{T}$ is a complete level-k type structure for μ if it is a level-k type structure that satisfies the following additional requirement: For each belief that assigns probability 1 to the m-types, there is an $(m+1)$-type of the player that holds that belief.

We can always find a complete level-k type structure.

Proposition 4.3. *Fix an $\boldsymbol{\mu} = (\mu_i : i \in I)$. There exists a complete level-k type structure for $\boldsymbol{\mu}$, viz. $\mathcal{T}^*$, that satisfies the following property: If $\mathcal{T}$ is a countable level-k type structure for $\boldsymbol{\mu}$, then $\mathcal{T}^*$ induces the hierarchies of beliefs induced by $\mathcal{T}$.*

The proof of Proposition 4.3 constructs a particular level-k type structure $\mathcal{T}^* = (S_{-i}, T_i^*, \beta_i^* : i \in I)$. The construction has a rich set of 1-types, i.e., for each $\nu_i \in \Delta(S_{-i} \times T_{-i}^*)$ with $\text{marg}_{S_{-i}} \nu_i = \mu_i$, there is a 1-type in T_i^* that holds that belief.[5] Thus, there are no restrictions on the beliefs of 1-types aside from the requirement that their first-order beliefs coincide with the anchor. With this, condition (iii) implies that the construction has a rich set of 2-types. And so on.

That said, there are hierarchies of beliefs consistent with the anchor that cannot be induced by any level-k type structure, a fortiori any complete level-k type structure. See Example 8.1. (See, also, Section 8-A on strengthening Proposition 4.3.) Section 8-B discusses why this is immaterial from the perspective of the inference problem.

5 The Inference Problem

We will be interested in the case where the analyst observes the strategy played and wants to infer the maximum level of reasoning about rationality consistent with observed behavior.[6] Reasoning about rationality will be captured by the epistemic conditions of *rationality and* m^{th}*-order belief of rationality.*

5.1 Rationality and m^{th}-order Belief of Rationality

An epistemic game $(G, \mathcal{T})$ induces a set of **states** $S \times T$. So, a state describes a strategy-type pair for each player. Rationality and m^{th}-order belief of rationality is a property that a state may or may not possess.

Definition 5.1. Say (s_i, t_i) is **rational** if $s_i \in \mathbb{BR}_i[\text{marg}_{S_{-i}} \beta_i(t_i)]$.

So a strategy-type pair (s_i, t_i) is rational if s_i is a best response under the first-order belief associated with t_i, viz. $\text{marg}_{S_{-i}} \beta_i(t_i)$.

Definition 5.2. Say $t_i \in T_i$ **believes** $E_{-i} \subseteq S_{-i} \times T_{-i}$ if E_{-i} is Borel and $\beta_i(t_i)(E_{-i}) = 1$.

So a type t_i believes an event if it assigns probability 1 to the event (i.e., to the Borel set E_{-i}). Given some $E_{-i} \subseteq S_{-i} \times T_{-i}$, write

$$B_i(E_{-i}) = \{t_i \in T_i : \beta_i(t_i)(E_{-i}) = 1\}$$

for the set of types that believe E_{-i}. Note, if $E_{-i} = \emptyset$, then $B_i(E_{-i}) = \emptyset$.

Write R_i^1 for the set of rational strategy-type pairs. Inductively define R_i^m by

$$R_i^{m+1} = R_i^m \cap (S_i \times \mathrm{B}_i(R_{-i}^m)).$$

Set $R_i^\infty = \bigcap_{m \geq 1} R_i^m$.

Definition 5.3. The set of states at which there is **rationality and** m^{th}**-order belief of rationality** (**R**m**BR**) is $R^{m+1} = \prod_{i \in I} R_i^m$. The set of states at which there is **rationality and common belief of rationality** (**RCBR**) is $R^\infty = \prod_{i \in I} R_i^\infty$.

[5] There are alternate constructions of complete level-k type structures, which do not satisfy this richness property.

[6] Of course, at times, authors augment the dataset with other observed variables of interest. Our concern is what the analyst can learn from the observed play, which is the focus of many studies.

5.2 The Unrestricted Inference Problem

The unrestricted inference problem is not the focus of our interest. Nonetheless, it will serve as a useful benchmark to think about the restricted inference problem.

In the unrestricted inference problem, the analyst observes the strategy played. But the analyst does not observe the set of hierarchies of beliefs players consider possible, i.e., the relevant type structure $\mathcal{T}$. Nor is the analyst prepared to make a substantive assumption about those beliefs. So, the relevant inference question is: If the analyst observes s_i, what is the maximum m so that s_i is consistent with RmBR in some type structure. More informally, what is the maximum level of reasoning about rationality consistent with observed behavior?

The answer to this question will depend on whether (or not) the observed strategy is m-rationalizable: Set $S_i^0 = S_i$ and assume the sets S_i^m have been defined. A strategy s_i is in S_i^{m+1} if and only if there exists some $\nu_i \in \Delta(S_{-i})$ with: (i) $s_i \in \mathbb{BR}_i[\nu_i]$, and (ii) $\nu_i(S_{-i}^m) = 1$. The set S_i^m is the set of **m-rationalizable strategies** for player i. The set $S_i^\infty = \bigcap_{m\geq 1} S_i^m$ is the set of **rationalizable strategies** for player i.

Proposition 5.1 (Known Result)**.** *Fix an epistemic game $(G, \mathcal{T})$.*

(i) For each $m \geq 1$, $\text{proj}_S R^m \subseteq S^m$.

(ii) If $\mathcal{T}$ is type-complete, for each $m \geq 1$, $\text{proj}_S R^m = S^m$.

(iii) If $\mathcal{T}$ is type-complete with compact type sets and continuous belief maps, $\text{proj}_S R^\infty = S^\infty$.

See Brandenburger and Dekel (1987), Tan and Werlang (1988), Battigalli and Siniscalchi (2002), and Friedenberg and Keisler (2021) for versions of this known result.

To understand how the result speaks to the unrestricted inference problem, consider two cases. First, suppose the analyst observes $s_i \in S_i^m \backslash S_i^{m+1}$, i.e., the analyst observes the player choose a strategy that is m- but not $(m+1)$-rationalizable. Then the analyst concludes the behavior is consistent with, at most, R$(m-1)$BR, i.e., m rounds of reasoning about rationality. In particular, s_i is consistent with R$(m-1)$BR in a type-complete type structure (part (ii)), but is inconsistent with RmBR in any other structure (part (i)).

Second, suppose the analyst observes $s_i \in S_i^\infty$. Then, there is a type-complete structure so that, for each m, the is a type t_i is consistent with RCBR (part (iii)). In this sense, s_i is consistent with unbounded reasoning about rationality.

5.3 The Restricted Inference Problem

In the restricted inference problem, the analyst is prepared to make the substantive assumption that hierarchies of beliefs are generated by some anchor $\boldsymbol{\mu}$. Thus, the relavent inference question is: If the analyst observes s_i, what is the maximum m so that s_i is consistent with RmBR in some level-k type structure for $\boldsymbol{\mu}$. One might think that the answer is tied to the level-k solution concept (for $\boldsymbol{\mu}$). However, as the next result indicates, it is not:

Theorem 5.1. *Fix an epistemic game $(G, \mathcal{T})$, where $\mathcal{T}$ is a level-k type structure for $\boldsymbol{\mu}$.*

(i) For each $m \geq 1$, $\text{proj}_S R^m \subseteq S^m$.

(ii) If $\mathcal{T}$ is a complete level-k type structure for $\boldsymbol{\mu}$, for each $m \geq 1$, $\text{proj}_S R^m = S^m$.

So, despite the fact that the analyst makes the substantive assumption that the hierarchies of beliefs are generated by a particular anchor $\boldsymbol{\mu}$, the nature of the inference problem is similar to the unrestricted inference problem: If the analyst observes a strategy that is m- but not $(m+1)$-rationalizable, then the analyst concludes the behavior is consistent with, at most, R$(m-1)$BR in any level-k type structure for $\boldsymbol{\mu}$. In particular, s_i is consistent with R$(m-1)$BR in a complete level-k type structure for $\boldsymbol{\mu}$ (part (ii)) but is inconsistent with RmBR is any level-k type structure for $\boldsymbol{\mu}$ (part (i)).

Note, if the analyst observes $s_i \in S_i^\infty$, then the conclusion is more subtle. Part (ii) says that the analyst cannot put a bound on reasoning about rationality, in the following sense: In a complete level-k type structure for $\boldsymbol{\mu}$, the strategy s_i is consistent with RmBR for each m. That is, in a complete level-k type structure, there are types $t_i^1, t_i^2, \ldots$ so that, for each m, $(s_i, t_i^m) \in R_i^m$. (Note, in general, t_i^m will *not* be an m-type.) However, this stops short of saying that s_i is consistent with RCBR. In fact, it may not be consistent with RCBR, as the following example indicates.

5.4 Proof of Theorem 5.1

We now turn to prove Theorem 5.1. Part (i) is an implication of Proposition 5.1's part (i). For part (ii) it suffices to show the reverse inclusion. In particular, we show the following: If $s_i \in S_i^m$, then there exists a $(m+1)$-type $t_i^{m+1} \in T_i^{m+1}$ so that $(s_i, t_i^{m+1}) \in R_i^m$. The proof is by induction on m.

First, fix $s_i \in S_i^1$. Then there exists some $\nu_i \in \Delta(S_{-i})$ such that s_i is a best response under ν_i. There exists $t_i^2 \in T_i^2$ such that $\text{marg}_{S_{-i}} \beta_i(t_i^2) = \nu_i$. As such, $(s_i, t_i^2) \in R_i^1$.

Next, assume the result holds for m. Fix $s_i \in S_i^{m+1}$. Then there exists some $\nu_i \in \Delta(S_{-i})$ such that s_i is a best response under ν_i and $\nu_i(S_{-i}^m) = 1$. By the induction hypothesis, there is a mapping $f_{-i}^m : S_{-i}^m \to T_{-i}^{m+1}$ such that $(s_{-i}, f_{-i}^m(s_{-i})) \in R_{-i}^m$. Construct $\hat{\nu}_i \in \Delta(S_{-i} \times T_{-i})$ so that $\hat{\nu}_i(s_{-i}, f_{-i}^m(s_{-i})) = \nu_i(s_{-i})$. In a complete level-$k$ type structure, there exists some $t_i^{m+2} \in T_i^{m+2}$ such that $\beta_i(t_i^{m+2}) = \hat{\nu}_i$. Since $\text{marg}_{S_{-i}} \beta_i(t_i^{m+2}) = \nu_i$, $(s_i, t_i^{m+2}) \in R_i^1$. Moreover, for each $n \leq m$, R_{-i}^n is Borel (Lemma C.3) and $\text{Supp}\, \beta_i(t_i^{m+2}) \subseteq R_{-i}^m \subseteq R_{-i}^n$. So, t_i^{m+1} believes R_{-i}^n for each $n \leq m$. As such, $(s_i, t_i^{m+2}) \in R_i^{m+1}$.

6 The Level-k Inference Problem

Theorem 5.1 raises the question: If we identify a subject as level-m but not level-n for $n > m$, what can we infer about the nature of the subject's reasoning. To address the question, we begin by providing an epistemic characterization of the level-k solution concept. We then discuss what the characterization means from the perspective of inferring reasoning about rationality.

6.1 The Level-k Solution Concept

Often, papers define the Level-k concept relative to a specific game. Because we want to define the concept for all (simultaneous-move) games, we introduce an abstract definition. We then discuss choices made in adopting the definition.

Definition 6.1. Set $L_i^1 = \mathbb{BR}_i[\mu_i]$. Assume the sets L_i^m have been defined. Let L_i^{m+1} be the set of strategies s_i so that there exists some $\nu_i \in \Delta(S_{-i})$ satisfying

(i) $s_i \in \mathbb{BR}_i[\nu_i]$, and

(ii) $\nu_i(L^m_{-i}) = 1$.

Say a strategy is **level-m (for $\boldsymbol{\mu}$)** if $s_i \in L^m_i$. Call the set L^m_i as **i's level-m behavior (for $\boldsymbol{\mu}$)** and call the set $L^m = \prod_{i \in I} L^m_i$ the **level-m behavior (for $\boldsymbol{\mu}$)**. The **level-k solution concept (for $\boldsymbol{\mu}$)** is the profile $(L^1, L^2, \ldots)$.

The level-k solution concept exogenously fixes a profile of first-order beliefs $\boldsymbol{\mu} = (\mu_i : i \in I)$, where μ_i reflects i's beliefs about the strategies others play. It then iterates best responses relative to those beliefs. Level-1 behavior is the set of strategy profiles $(s_i : i \in I)$ where each s_i is a best response under i's anchor. Level-2 behavior is the set of strategy profiles $(s_i : i \in I)$ where each s_i is a best response under a belief that assigns probability one to the level-1 behavior of other players. Etc.

Remark 6.1. Our definition allows for the fact that the sets L^m_i may not be a singleton. In fact, there are prominent examples where the level-k solution concept has been applied, despite the fact that there are multiple best responses. For instance, consider a 3-player beauty contest game (Nagel, 1995), where players simultaneously choose a number in $\{1, 2, 3, 4, 5\}$. A player wins if their choice is closest to $\frac{2}{3}$ of the average; they loose if some other bid is closer to $\frac{2}{3}$ of the average. Ties split the win equally. If the anchor assigns probability 1 to the arithmatic mean 3, then bidding either of 1 or 2 is a best response.

When there are multiple best responses, some papers assume players have a uniform belief over those best responses. So, in the beauty contest example of the previous paragraph, a level-2 strategy must be a best response under a belief that assigns $\frac{1}{2} : \frac{1}{2}$ to $1 : 2$. This imposes a secondary exogenous restriction on beliefs—but one that depends on iterative best responses. We discuss this further in Section 8-D.

6.2 Epistemic Foundations for Level-k

Theorem 6.1. *Fix an epistemic game $(G, \mathcal{T})$, where $\mathcal{T}$ is a level-k type structure for $\boldsymbol{\mu}$. For each player i, fix covers $\mathcal{C}_i = \{T^m_i : m = 1, 2, \ldots\}$ satisfying conditions (i)-(ii) of Definition 4.1 (resp. (i)-(ii)-(iii) of Definition 4.2, if $\mathcal{T}$ is a complete level-k type structure).*

(i) For each $m \geq 1$, $\mathrm{proj}_{S_i}(R^m_i \cap (S_i \times T^m_i)) \subseteq L^m_i$.

(ii) For each $m \geq 1$, if $\mathcal{T}$ is a complete level-k type structure for $\boldsymbol{\mu}$, $\mathrm{proj}_{S_i}(R^m_i \cap (S_i \times T^m_i)) = L^m_i$.

Much like Theorem 5.1, Theorem 6.1 fixes a level-k type structure for $\boldsymbol{\mu}$. Refer to Figure TBA. Whereas Theorem 5.1 focused on the behavioral implications of R$(m-1)$BR, Theorem 6.1 focuses on the behavioral implications of R$(m-1)$BR for *only the m-types*. Part (i) says that, if the m-types engage in R$(m-1)$BR, their behavior is level-m (for $\boldsymbol{\mu}$). Part (ii) adds that, in any complete level-k type structure, any level-m strategy for $\boldsymbol{\mu}$ is consistent with R$(m-1)$BR *for an m-type.*

To better understand the Theorem, fix a level-k type structure for $\boldsymbol{\mu}$ (not necessarily a complete level-k type structure). A strategy is level-1 strategy for $\boldsymbol{\mu}$ if and only if there is a 1-type t_i so that (s_i, t_i) is rational. (See Proposition C.1 part (i).) Note, this conclusion is stronger than that in part (i) and must only hold for $m = 1$. In particular, a strategy s_i may be level-2 for $\boldsymbol{\mu}$ even if there is no 2-type t_i so that (s_i, t_i) is consistent with R1BR. The next example illustrates this claim.

Example 6.1. Refer to the game in Figure 6.1. Consider an anchor $\boldsymbol{\mu} = (\mu_1, \mu_2)$ with each $\mu_i(D_{-i}) = 1$. Observe that

$$L^m_i = \{U_i, M_i\} = S^m_i$$

for each $m \geq 1$. We next show that there is a level-k type structure for $\boldsymbol{\mu}$ so that (i) each m-rationalizable strategy is consistent with RmBR, but (ii) there is a level-k strategy for $\boldsymbol{\mu}$ so that some $s_i \in L_i^2$ is inconsistent with R1BR for every 2-type.

	U_{-i}	M_{-i}	D_{-i}
U_i	1	0	1
M_i	0	1	1
D_i	-1	-1	-1

Figure 6.1

Next define a level-k type structure for $\boldsymbol{\mu}$. Set each $T_i = \{t_i, v_i\} \times \mathbb{N}$. The belief maps are defined as follows: First, $\beta_i(t_i, 1)(D_{-i}, (t_{-i}, 2)) = \beta_i(v_i, 1)(D_{-i}, (v_{-i}, 2)) = 1$ and $\beta_i(t_i, 2)(U_{-i}, (t_{-i}, 1)) = \beta_i(v_i, 2)(U_{-i}, (v_{-i}, 1)) = 1$. Second, $\beta_i(t_i, 3)(U_{-i}, (t_{-i}, 2)) = 1$ but $\beta_i(v_i, 3)(D_{-i}, (v_{-i}, 2)) = 1$. Third, for each $n \geq 4$, $\beta_i(t_i, n)(U_{-i}, (t_{-i}, n-1)) = \beta_i(v_i, n)(M_{-i}, (v_{-i}, n-1)) = 1$. Note, this is a level-$k$ type structure for $\boldsymbol{\mu}$ associated with covers $\mathcal{C}_i = \{\{t_i, v_i\} \times \{m\} : m \geq 1\}$.

For each m, $\text{proj}_{S_i} R_i^m = \{U_i, M_i\}$. However, $\text{proj}_{S_i}(R_i^2 \cap T_i^2) = \{U_i\} \subsetneq L_i^m$.

Example 6.1 features a "rich" level-k type structure, in the sense that there are enough beliefs so that all the m-rationalizable strategies are consistent with R$(m-1)$BR. Thus, for this specific type structure, part (i) of Theorem 5.1 can be strengthened from inclusion to equality. Despite the type structure being rich in this sense, it does not have a rich set of 2-types. As a consequence, there are level-2 strategies that are inconsistent with R1BR for each 2-type. Part (ii) of Theorem 6.1 says implies that, when there is a "rich" set of 2-types (in the sense of the requirement associated with a complete level-k type structure), any level-2 strategies is consistent with R1BR for some 2-type.

While a complete level-k type structure features a sufficiently "rich" set of 2-types, 3-types, etc., it is important to note that it does not induce a rich set of beliefs: In particular, we saw that a complete level-k type structure *cannot* induce all hierarchies of beliefs. Moreover, a complete level-k type structure *cannot* be type-complete. See Propositions 4.1-4.2.

6.3 Identifying Levels of Reasoning about Rationality

Suppose the analyst observes a player choose some strategy s_i^* so that (i) s_i^* is level-m ($m \geq 1$) for $\boldsymbol{\mu}$, but (ii) s_i^* is not level-n (for $\boldsymbol{\mu}$) for any $n > m$. What can the analyst infer about how the player reasons about rationality? We first address the question in the context of the unrestricted inference problem, then in the context of the restricted inference problem, and finally in the context of Theorem 6.1. To do so, we use the following fact: If s_i^* is level-m then s_i^* is m-rationalizable. (See Lemma C.4.)

In the unrestricted inference problem, the analyst only observes the strategy s_i^* and the analyst is not prepared to make an assumption about the hierarchies of beliefs that players consider possible. Since s_i^* is level-m, the analyst concludes that s_i^* is consistent with R$(m-1)$BR in some type structure. Because an $(m+1)$-rationalizable strategy need not be level-$(m+1)$, the strategy s_i^* might well be consistent with RmBR in some type structure, even though it is not level-$(m+1)$ for $\boldsymbol{\mu}$. The analyst can only conclude that s_i^* is inconsistent with RmBR if the strategy is not $(m+1)$-rationalizable. This is an implication of Proposition 5.1.

In the restricted inference problem, the analyst is willing to make a substantive assumption about the players' beliefs:

Identifying Assumption 1. *There is some anchor $\boldsymbol{\mu}$ and some $n \geq 1$ so that the player who chose s_i^* has the beliefs associated with an n-type in a level-k type structure for $\boldsymbol{\mu}$.*

Theorem 5.1 implies that, despite this identification assumption, the nature of the inference does not change: The analyst can conclude that s_i^* is consistent with R$(m-1)$BR, but cannot rule out that it is also consistent with RmBR, unless s_i^* is also fails $(m+1)$-rationalizability.

Theorem 6.1 suggests a stronger conclusion, based on an additional auxiliary assumption above Assumption 1:

Identifying Assumption 2. *If a player is an n-type in some level-k type structure for $\boldsymbol{\mu}$, then the player reasons according to R$(n-1)$BR.*

Under Assumptions 1-2, the analyst can conclude that s_i^* is consistent with R$(m-1)$BR and inconsistent with RnBR for any $n \geq m$: Since, for each $n \geq m+1$, s_i^* is not level-n (for $\boldsymbol{\mu}$), there is no level-k type structure (for $\boldsymbol{\mu}$) and n-type thereof t_i, so that (s_i^*, t_i) is consistent with R$(n-1)$BR. (This uses Theorem 6.1.) Then, the identifying assumptions rule out that the behavior s_i^* was generated by a player that reasons according to RmBR, a fortiori RnBR for any $n > m$.

It is worth emphasizing the nature of this approach to identification, especially relative to standard critiques in the literature. It is understood that the level-k approach implicitly assumes that behavior is generated by subjects who have (partial) beliefs (of some order) induced by an anchor. This assumption fits with Assumption 1 and has itself received criticism. (Refer back to page 4.) The analysis here highlights the importance of Assumption 2, above and beyond Assumption 1. A generous interpretation of Assumption 2 is: If subjects hold certain partial n^{th}-order beliefs, then there reason according to R$(n-1)$BR.[7] Even under this interpretation, it requires an assumption that particular n^{th}-order beliefs determine how a player reasons about rationality—an assumption that would be hard to verify (or falsify) in practice.

Remark 6.2. [**TODO:** Add discussion about the difference in the conclusion on reasoning about rationality vs reasoning]

7 Applications

[**TODO:** Insert applications]

8 Discussion

A. Complete Level-k Type Structures and Hierarchies Consistent with the Anchor One might conjecture that a complete level-k type structure for $\boldsymbol{\mu}$ induces all hierarchies of beliefs consistent with the anchor. However, this is not the case. We begin with an example.

Example 8.1. Consider a two-player game where each $S_i = \{\Box_i, \Diamond_i\}$. For each player i, there is a hierarchy of beliefs $h_{i,\Box} = (h^1_{i,\Box}, h^2_{i,\Box}, \ldots)$ where it is commonly believed that the other player chooses $\Box_{-i}$: So, $h^1_{i,\Box}(\Box_{-i}) = 1$ and $h^{m+1}_{i,\Box}(\Box_{-i}, \ldots, h^m_{-i,\Box}) = 1$. Also, for each player i, there is a hierarchy of beliefs $h_i = (h^1_i, h^2_i, \ldots)$ with $h^1_i(\Box_{-i}) = \frac{2}{3}$, $h^{m+1}_i(\Box_{-i}, \ldots, h^m_{-i,\Box}) = \frac{2}{3}$, and $h^{m+1}_i(\Diamond_{-i}, \ldots, h^m_{-i,\Box}) = \frac{1}{3}$. (So, h^2_i

[7] This is indeed generous. In particular, n-types are associated with certain partial n^{th}-order beliefs, but the partial n^{th}-order beliefs do not uniquely determine whether a type is an n-type.

assigns probability $\frac{2}{3}$s to "the other player plays $\square_{-i}$ and believes I play $\square_i$" and probability $\frac{1}{3}$s to "the other player plays $\diamond_{-i}$ and believes I play $\square_i$.")

Now consider an anchor $\boldsymbol{\mu} = (\mu_1, \mu_2)$ where, for each i, $\mu_i(\square_{-i}) = \frac{2}{3}$. Note, that h_i is a hierarchy of beliefs consistent with the anchor, since $h_i^1 = p_{i,\boldsymbol{\mu}}^1$. However, there is no level-k type structure (including a complete level-k type structure) that induces the hierarchy $h_i = (h_i^1, h_i^2, \ldots)$. We give the intuition why here and complete the proof in Appendix D.

Fix a level-k type structure for $\boldsymbol{\mu} = (\mu_i : i \in I)$ and, for each $i \in I$, let $\mathcal{C}_i = \{T_i^m : m = 1, 2, \ldots\}$ be a Borel cover so that $(\mathcal{C}_i : i \in I)$ jointly satisfy conditions (i)-(ii) of Definition 4.1. Suppose, contra hypothesis, there exists some type $t_i \in T_i$ with $\delta_i(t_i) = h_i$. Then, there must exist some type $t_{-i,\square} \in T_{-i}$ with $\delta_{-i}(t_{-i,\square}) = h_{-i,\square}$. (See Lemma D.1.) But, there is no such type $t_{-i,\square} \in T_{-i}$. (See Lemma D.2.) Intuitively: The 1-types have first-order beliefs distinct from $h_{i,\square}^1$. Since the 2-types must assign probability one to 1-types, this implies that the 2-types have second-order beliefs distinct from $h_{i,\square}^1$. And so on.

The example points to a more general phenomena. A level-k type structure (a fortiori, a complete level-k type structures) cannot induce hierarchies of beliefs where the first-order beliefs coincide with the anchor and higher-order beliefs assigns positive probability to beliefs that are inconsistent with the anchor. As a consequence, it also cannot induce hierarchies of beliefs that assign positive probability to such hierarchies. Etc. Put differently, level-k type structures (a fortiori, complete level-k type structures) impose the substantive requirement: Not only are players beliefs consistent with the anchor, they believe other players' beliefs are consistent with the anchor, etc.[8] As a consequence:

Proposition 8.1. *Fix a non-degenerate anchor $\boldsymbol{\mu}$, i.e., an anchor where no player assigns probability 1 to a strategy profile. If $\mathcal{T}$ is a level-k type structure for $\boldsymbol{\mu}$, then $\mathcal{T}$ does not induce all hierarchies of beliefs consistent with $\boldsymbol{\mu}$.*

One might instead hope for the following: If a hierarchy can be induced by a level-k type structure for $\boldsymbol{\mu}$, then any complete level-k type structure must also induce that hierarchy. However, a close inspection of Definition 4.2 indicates why this need not be the case. While a complete level-k type structure requires a rich set of 2-types, 3-types, etc., it does not require a rich set of 1-types.

The proof of Proposition 4.3 constructs a particular complete level-k type structure $\mathcal{T}^* = (T_i^*, \beta_i^* : i \in I)$ that does have a rich set of 1-types: For every belief $\nu_i \in \Delta(S_{-i} \times T_{-i}^*)$ with $\text{marg}_{S_{-i}} \nu_i = \mu_i$, there is a 1-type in T_i^* with $\beta_i^*(t_i^*) = \nu_i$. For this reason, any hierarchy of beliefs that can be induced by a countable level-k type structure can be induced by the constructed complete level-k type structure. Appendix B.4 discusses the technical difficulty in extending the result to any level-k type structure.

B. Complete Level-k Type Structures and Inference We saw that a complete level-k type structure need not induce all hierarchies of beliefs consistent with the anchor. Despite this, from the perspective of inferring the level of reasoning about rationality, it suffices to focus on level-k and complete level-k type structures. To understand why, recall that in any type structure, the set of strategies consistent with R$(m-1)$BR must be contained in the m-rationalizable strategies. (Refer to Proposition 5.1(i).) The same holds if we replace "any type structure" with "any hierarchy structure" (i.e., any belief-closed subset—or even any subset—of hierarchies of beliefs). Since any m-rationalizable strategy is consistent

[8] Of course, one might want to implose this substantive requirement. The literature is, arguably, silent on whether this is desired.

with R$(m-1)$BR strategy in a complete level-k type structure (Theorem 5.1(ii)), the focus on complete level-k type structures is without loss of inference.

C. Definition of Level-k Type Structures A level-k type structure (Definition 4.1) requires that, for each player i, we find a cover that satisfies two properties. It does not require that the associated covers be unique. Indeed, they may not be; see Example D.2. It also does not require that the cover is a partition. Indeed, they may not be; see Example D.1.

A complete level-k type structure (Definition 4.2) is associated with covers that satisfy three criteria. While these covers need not be a partition, the construction of a complete level-k type structure in Proposition 4.3 does involve partitional covers. We do not know if adding a partitional requirement imposes substantive assumptions.

D. Definition of the Level-k Solution Concept Definition 6.1 allows for the fact that there may be multiple best responses to a given distribution on strategies. This is not simply a theoretical possibility but a feature of important level-k analyses. As pointed out in Remark 6.1, some papers instead assume that players have a uniform belief about best responses. This imposes a secondary exogenous restriction on beliefs—but one that depends on iterative best responses. This additional restriction only serves to reinforce the message of the paper: It might suggest lower levels of reasoning about rationality than is consistent with the data, since it may suggest that the level-k bound is lower than that suggested by Definition 6.1.

Theorem 6.1 can be seen as providing foundations for this level-k solution concept, as specified by Definition 6.1. From the perspective of foundations, it is important that we focus on this generalized level-k solution concept. The epistemic approach takes, as given, the set of hierarchies of beliefs players consider possible (i.e., a type structure); it then goes on to impose epistemic conditions relative to those hierarchies (i.e., RmBR is applied relative to a type structure). The restriction to a uniform belief over best responses proceeds in a different direction: It derives first-order beliefs based on best responses (to other beliefs).

E. Foundations for Level-k Theorem 6.1 provides epistemic foundations for the level-k solution concept. These foundations are quite different from foundations for other solution concepts: The foundations rest on associating different hierarchies of partial beliefs with different epistemic conditions. In doing so, it allows the researcher to make different epistemic assumptions (i.e., R1BR, R2BR, etc...) based on different hierarchies of partial beliefs. By contrast, the typical approach (in epistemic game theory) will simply say whether a hierarchy of beliefs is or is not consistent with a particular epistemic assumption.

The foundations are cast in a typical epistemic framework, where types are associated with hierarchies of beliefs. This approach describes players as actors that do not face limitations on their ability to engage in interactive reasoning—i.e., their ability to specify all sentences of the form "I think that you think that I think ...." However, often, the level-k solution concept is motivated by a stipulation that players have a limited ability to engage in such sentences. On the one hand, Theorem 6.1 indicates that this stipulation is not needed—that the level-k solution concept does not *require* limits on the ability to engage in interactive reasoning. On the other hand, one might wonder if the foundations hinge on unlimited ability to engage in interactive reasoning. They do not: We can recast the analysis here in terms of an epistemic model where epistemic types only induce finite-order beliefs (as in Heifetz and Kets, 2018 or Kets, 2010). The key is that the epistemic conditions of RmBR depend only on the $(m+1)^{\text{th}}$-order beliefs. See Appendix D.

F. RmBR Behavior of k-Types: Dominated Anchors [**NOTE:** We had the lemma below. Do we find it useful? Should we keep it?]

Lemma 8.1. *Fix an epistemic game $(G, \mathcal{T})$, where $\mathcal{T}$ that is a level-k type structure for $\boldsymbol{\mu}$. If, for each i, $\mu_i(S_{-i} \backslash S^1_{-i}) > 0$, then*

$$\bigcup_{k \geq m} \operatorname{proj}_{S_i} (R^m_i \cap (S_i \times T^m_i)) = \operatorname{proj}_{S_i} R^m_i.$$

In the specific case where the anchor assigns positive probability to a dominated strategy, the R$(m-1)$BR behavior coincides with the R$(m-1)$ behavior of $k \geq m$ types. This arises because, for such an anchor, there are no $k \leq m-1$ types consistent with R$(m-1)$BR. (See Lemma D.3.) In particular, 1-types assign positive probability to irrational strategy-type pairs; as such, they are inconsistent with R1BR. With this, 2-types assign probability 1 to strategy-type pairs inconsistent with R1BR; as such, they are inconsistent with R2BR. And so on.

Appendix A Mathematical Preliminaries

It will be useful to begin with some mathematical preliminary.

Lemma A.1. *Let Ω_1, Ω_2 be metrizable spaces. Then the mapping $\operatorname{marg}_{\Omega_1} : \Delta(\Omega_1 \times \Omega_2) \to \Delta(\Omega_1)$ is continuous.*

Proof. Let $\operatorname{proj} : \Omega_1 \times \Omega_2 \to \Omega_1$ be the projection mapping, i.e., mapping $\operatorname{proj}(\omega_1, \omega_2) = \omega_1$. Note that proj is continuous: If $U_1 \subseteq \Omega_1$ is open, then $(\operatorname{proj})^{-1}(U_1) = U_1 \times \Omega_2$ is open. Thus, $\underline{\operatorname{proj}} : \Delta(\Omega_1 \times \Omega_2) \to \Delta(\Omega_1)$ is continuous. (See Theorem 15.14 in Aliprantis and Border (2007).) Finally, observe that $\underline{\operatorname{proj}}(\nu) = \operatorname{marg}_{\Omega_1}\nu$. (For each event $E_1 \subseteq \Omega_1$, $\operatorname{marg}_{\Omega_1}\nu(E_1) = \nu(E_1 \times \Omega_2) = \underline{\operatorname{proj}}(\nu)(E_1)$.) ■

Appendix B Proofs for Sections 3-4

B.1 Type Structures Induce Hierarchies of Beliefs

Fix a type structure $\mathcal{T} = (S_{-i}, T_i, \beta_i : i \in I)$. We will inductively define measurable maps $\rho^m_i : S_{-i} \times T_{-i} \to X^m_i$ and $\delta^m_i : T_i \to H^m_i$. First, set $\rho^1_i = \operatorname{proj}_{S_{-i}}$ and $\delta^1_i = \underline{\rho}^1_i \circ \beta_i$. Note, ρ^1_i is measurable and so $\underline{\rho}^1_i$ is measurable. From this and the fact that β_i is measurable, δ^1_i is measurable.

Now, assume the measurable maps $\rho^m_i : S_{-i} \times T_{-i} \to X^m_i$ and $\delta^m_i : T_i \to H^m_i$ have been defined. Set

$$\rho^{m+1}_i(s_{-i}, t_{-i}) = (\rho^m_i(s_{-i}, t_{-i}), \delta^m_{-i}(t_{-i})).$$

Note, since ρ^m_i and δ^m_{-i} are measurable, so is ρ^{m+1}_i. Then set $\delta^{m+1}_i = \underline{\rho}^{m+1}_i \circ \beta_i$. Since ρ^{m+1}_i is measurable and so $\underline{\rho}^{m+1}_i$ is measurable. From this and the fact that β_i is measurable, δ^{m+1}_i is measurable.

The following standard Lemmata will be of use.

Lemma B.1. *For each $t_i \in T_i, \delta^1_i(t_i) = \operatorname{marg}_{S_{-i}}\beta_i(t_i)$.*

Proof. Fix some $s_{-i} \in S_{-i}$. Note,

$$\delta_i^1(t_i)(s_{-i}) = \beta_i(t_i)((\rho_i^1)^{-1}(\{s_{-i}\})) = \beta_i(t_i)(S_{-i} \times T_{-i}),$$

as desired. ∎

Lemma B.2. *If $\delta_i^m(t_i) = h_i^m$ then, for each $n \le m$, $\delta_i^n(t_i) = \operatorname{marg}_{X_{-i}^n} h_i^n$,*

Proof. Insert proof. ∎

Lemma B.3. *For each $m \ge 1$, $\rho_i^{m+1}(s_{-i}, t_{-i}) = (s_{-i}, \delta_{-i}^1(t_{-i}), \ldots, \delta_{-i}^m(t_{-i}))$.*

Proof. For $m = 1$, this is immediate. Assume the statement is true for $m \ge 2$, so that $\rho_i^{m+1}(s_{-i}, t_{-i}) = (s_{-i}, \delta_{-i}^1(t_{-i}), \ldots, \delta_{-i}^m(t_{-i}))$. Then, $\rho_i^{m+2}(s_{-i}, t_{-i}) = (s_{-i}, \delta_{-i}^1(t_{-i}), \ldots, \delta_{-i}^m(t_{-i}), \delta_{-i}^{m+1}(t_{-i}))$, as desired. ∎

B.2 Proof of Proposition 4.1

Fix a level-k type structure for $\boldsymbol{\mu} = (\mu_i : i \in I)$ and, for each $i \in I$, let $\mathcal{C}_i = \{T_i^m : m = 1, 2, \ldots\}$ be a Borel cover so that $(\mathcal{C}_i : i \in I)$ jointly satisfy conditions (i)-(ii) of Definition 4.1. The following Lemma will establish Proposition 4.1.

Lemma B.4. *For each $m \ge 1$, $\eta_i^m(\delta_i^m(T_i^m)) \subseteq \{p_{i,\boldsymbol{\mu}}^m\}$.*

Proof. The case of $m = 1$ is immediate. Assume the claim holds for $m \ge 2$. Fix some $t_i \in T_i^{m+1}$ and write $h_i^{m+1} = \delta_i^{m+1}(t_i)$. We will show that $\eta_i^{m+1}(h_i^{m+1}) = p_{i,\boldsymbol{\mu}}^{m+1}$.

Fix

$$E_i^{m+1} = X_i^m \times \prod_{j \ne i} (\eta_j^m)^{-1}(\{p_{j,\boldsymbol{\mu}}^m\})$$

and note that $E_i^{m+1} \subseteq X_i^{m+1}$. Note that

$$S_{-i} \times T_{-i}^m \subseteq (\rho_i^{m+1})^{-1}(E_i^{m+1}).$$

To see this, fix $(s_{-i}, t_{-i}) = (s_j, t_j : j \ne i) \in S_{-i} \times T_{-i}^m$. By the induction hypothesis, $\eta_j^m(\delta_j^m(t_j^m)) = p_{j,\boldsymbol{\mu}}^m$. Thus, $\rho_i^{m+1}(s_{-i}, t_{-i}) \in E_i^{m+1}$, as stated.

Now observe that E_i^{m+1} is measurable set, since each η_j^m is measurable. Thus,

$$h_i^{m+1}(E_i^{m+1}) = \beta_i(t_i)((\rho_i^{m+1})^{-1}(E_i^{m+1})) \ge \beta_i(t_i)(S_{-i} \times T_{-i}^m) = 1$$

and so $h_i^{m+1}(E_i^{m+1}) = 1$. We use this fact to show that $\eta_i^{m+1}(h_i^{m+1}) = p_{i,\boldsymbol{\mu}}^{m+1}$. In particular, suppose $\eta_i^{m+1}(h_i^{m+1}) = \tilde{p}$. Note,

$$\tilde{p}(\{p_{-i,\boldsymbol{\mu}}^m\}) = h_i^{m+1}((\hat{\eta}_{-i}^m)^{-1}(\{p_{-i,\boldsymbol{\mu}}^m\})) = h_i^{m+1}(E_i^{m+1}) = 1.$$

Thus, $\tilde{p} = p_{i,\boldsymbol{\mu}}^{m+1}$ as desired. [**NOTE:** check that it is indeed obvious that $(\hat{\eta}_{-i}^m)^{-1}(\{p_{-i,\boldsymbol{\mu}}^m\}) = E_i^{m+1}$.] ∎

B.3 Proof of Proposition 4.2

Proof of Proposition 4.2. Fix a level-k type structure for $\boldsymbol{\mu}$, viz. $\mathcal{T} = (S_{-i}, T_i, \beta_i : i \in I)$. For each $i \in I$, write $\mathcal{C}_i = \{T_i^1, T_i^2, \ldots\}$ for Borel covers of T_i that jointly satisfy conditions conditions (i)-(ii) of

Definition 4.1.

Suppose, contra hypothesis, $\mathcal{T}$ is type-complete. First, let $\nu_i \in \Delta(S_{-i} \times T_{-i})$ where (i) $\text{marg}_{S_{-i}} \nu_i = \mu_i$ and (ii) for each $m \geq 1$, $\nu_i(S_{-i} \times T_{-i}^m) > 0$. Since $\mathcal{T}$ is type-complete, there exists a type $t_i^1 \in T_i$ with $\beta_i(t_i^1) = \nu_i$. Using the fact that $\beta_i(t_i^1)(S_{-i} \times T_{-i}^m) > 0$ for each m, it follows that $t_i^1 \notin T_i^m$ for all $m \geq 2$. Thus, $t_i^1 \in T_i^m$ if and only if $m = 1$.

Now, observe that there is a $\varphi_i \in \Delta(S_{-i} \times T_{-i})$ so that (i) $\text{marg}_{S_{-i}} \varphi_i \neq \mu_i$, and (ii) $\varphi_i(S_{-i} \times \{t_{-i}^1\}) = 1$. Since $\mathcal{T}$ is type-complete, there exists a type $t_i^2 \in T_i$ with $\beta_i(t_i^2) = \varphi_i$. Given that $\text{marg}_{S_{-i}} \varphi_i \neq \mu_i$, $t_i^2 \notin T_i^1$. Using the fact that $t_i^1 \in T_i^m$ if and only if $m = 1$, it follows that $t_i^2 \in T_i^m$ if and only if $m = 2$.

Finally, observe that there is a $\xi_i \in \Delta(S_{-i} \times T_{-i})$ so that (i) $\text{marg}_{S_{-i}} \xi_i \neq \mu_i$, (ii) $\xi_i(S_{-i} \times \{t_{-i}^1\}) > 0$, and (iii) $\xi_i(S_{-i} \times \{t_{-i}^2\}) > 0$. Since $\mathcal{T}$ is type-complete, there exists a type $t_i \in T_i$ with $\beta_i(t_i) = \xi_i$. Given that $\text{marg}_{S_{-i}} \xi_i \neq \mu_i$, $t_i \notin T_i^1$. Since $\xi_i(S_{-i} \times \{t_{-i}^1\}) > 0$ and $\xi_i(S_{-i} \times \{t_{-i}^2\}) > 0$ and there is no m with $t_{-i}^1, t_{-i}^2 \in T_i^m$, it follows that $t_i \notin T_i^m$ for each m. This contradicts the type structure being a level-k type structure for $\boldsymbol{\mu}$. ■

B.4 Proof of Proposition 4.3

B.4.0.1 Construction of a Complete Level-k Type Structure For each integer $m \geq 1$, let $T_i^{*,m} = [0,1] \times \{m\}$. Set $T_i^* = \bigcup_{m \geq 1} T_i^{*,m}$. Endow $T_i^{*,m}$ with a metric $d : T_i^* \times T_i^* \to \mathbb{R}$ so that $d((x_j, m_j), (x_\ell, m_\ell)) = \|x_j - x_\ell\|$ if $m_j = m_\ell$ and $d((x_j, m_j), (x_\ell, m_\ell)) = 2$ if $m_j \neq m_\ell$.

Lemma B.5. *Then (T_i^*, d) is a Polish space.*

Proof. Let $D_m = (\mathbb{Q} \cap [0,1]) \cap \{m\}$ and note that each D_m forms a countable dense subset of $[0,1] \times \{m\}$. Then set $D = \bigcup_{m \in \mathbb{Z}} (D_m \times \{m\})$. The set D is countable. It is also dense in T_i^*. (This follows from the fact that each open set in T_i^* must either be an open set in $[0,1] \times \{m\}$ or a union of such open sets.) Thus, (T_i^*, d) is separable.

Next observe that, for any Cauchy sequence $((x_j, m_j) : j = 1, 2, \ldots)$, there must be some J so that $m_j = m_J$ for all $j \geq J$. Thus, any Cauchy sequence converges and (T_i^*, d) is complete. ■

Lemma B.6.

(i) There exists an injective bimeasurable map $\chi_i^1 : T_i^{*,1} \to \Delta(S_{-i} \times T_{-i}^*)$ so that $\chi_i^1(T_i^{*,1}) = \{\nu_i \in \Delta(S_{-i} \times T_{-i}^*) : \text{marg}_{S_{-i}} \nu_i = \mu_i\}$.

(ii) For each $m \geq 2$, there exists an injective bimeasurable map $\chi_i^m : T_i^{*,m} \to \Delta(S_{-i} \times T_{-i}^*)$ so that $\chi_i^m(T_i^{*,m}) = \{\nu_i \in \Delta(S_{-i} \times T_{-i}^*) : \nu_i(S_{-i} \times T_{-i}^{*,m-1}) = 1\}$.

Proof. For part (i), begin by noting that both $T_i^{*,1} = [0,1] \times \{1\}$ and $\Delta(S_{-i} \times T_{-i}^*)$ are uncountable Polish spaces. (The latter follows from Lemma B.5.) Since $\{\nu_i \in \Delta(S_{-i} \times T_{-i}^*) : \text{marg}_{S_{-i}} \nu_i = \mu_i\}$ is a closed subset of $\Delta(S_{-i} \times T_{-i}^*)$, it too is Polish. (See Aliprantis and Border, 2007, pg. 74.) Moreover, $\{\nu_i \in \Delta(S_{-i} \times T_{-i}^*) : \text{marg}_{S_{-i}} \nu_i = \mu_i\}$ is uncountable. So, the claim follows from the Borel Isomorphism Theorem.

For part (ii), fix $m \geq 2$. Note that both $T_i^{*,m}$ and $\Delta(S_{-i} \times T_{-i}^{m-1,*})$ are uncountable Polish spaces. So, by the Borel Isomorphism Theorem, there exists a bimeasurable bijective map $\hat{\chi}_i^m : T_i^{*,m} \to \Delta(S_{-i} \times T_{-i}^{m-1,*})$. Also note that there exists an injective bimeasurable map $\hat{\psi}_i^m : \Delta(S_{-i} \times T_{-i}^{m-1,*}) \to \Delta(S_{-i} \times T_{-i}^*)$ so that

$$\hat{\psi}_i^m(\Delta(S_{-i} \times T_{-i}^{m-1,*})) = \{\nu_i \in \Delta(S_{-i} \times T_{-i}^*) : \nu_i(\Delta(S_{-i} \times T_{-i}^{*,m-1})) = 1\}.$$

Thus, $\hat{\psi}_i^m \circ \hat{\chi}_i^m$ is an injective bimeasurable map that satisfies the desired property. ■

For each i, let $\beta_i^* : T_i^* \to \Delta(S_{-i} \times T_{-i}^*)$ be defined so that $\beta_i(x,m) = \chi_i^m(x,m)$. Note, under this construction, β_i^* is not injective. But, if there exists $(x,m) \neq (x',m')$ but $\beta_i^*(x,m) = \beta_i^*(x',m')$, then either (i) $(x,m) \in [0,1] \times \{1\}$ and $(x',m') \notin [0,1] \times \{1\}$ or (ii) $(x',m') \in [0,1] \times \{1\}$ and $(x,m) \notin [0,1] \times \{1\}$.

Lemma B.7. *The map β_i^* is bimeasurable.*

Proof. Fix a Borel $E \subseteq S_{-i} \times T_{-i}^*$. Since each χ_i^m is measurable, each $(\chi_i^m)^{-i}(E)$ is Borel. Now observe that

$$(\beta_i^*)^{-1}(E) = \bigcup_{m \geq 1} (\chi_i^m)^{-i}(E)$$

is Borel. Thus, β_i^* is measurable.

Likewise, fix a Borel $E \subseteq T_i^*$. Since each χ_i^m is bimeasurable, each $\chi_i^m(E \cap T_i^{*,m})$ is Borel. From this

$$\beta_i^*(E) = \bigcup_{m \geq 1} \chi_i^m(E \cap T_i^{*,m})$$

is Borel. Thus, β_i^* is bimeasurable. ■ Note, this establishes that $(T_i^*, \beta_i^* : i \in I)$ is a type structure. Let $\rho_i^{*,m} : S_{-i} \times T_{-i}^* \to X_i^m$ (resp., $\delta_i^{*,m} : T_i^* \to H_i^m$) be the map from strategy-type pairs to the m^{th}-order space of uncertainty (resp. be the map from types to m^{th}-order beliefs).

Lemma B.8. *The type structure $(T_i^*, \beta_i^* : i \in I)$ is a complete level-k type structure.*

Proof. Observe that $\mathcal{C}_i^* = \{T_i^{*,m} : m = 1, 2, \ldots\}$ is a Borel cover that, by construction, satisfies conditions (i)-(ii)-(iii) of a complete level-k type structure. ■

B.4.0.2 Induces Hierarchies of Countable Level-k Type Structures

For the remainder of the argument, fix a level-k type structure $(T_i, \beta_i : i \in I)$. Then there exists a Borel cover $\mathcal{C}_i = \{T_i^m : m = 1, 2, \ldots\}$ that satisfies conditions (i)-(ii) of Definition 4.1. Let $\rho_i^m : S_{-i} \times T_{-i} \to X_i^m$ and $\delta_i^m : T_i \to H_i^m$ be the maps associated with this type structure.

Lemma B.9. *Suppose, for each i, T_i is countable. Then, for each m and each n, there is a map $f_i^{m,n} : T_i^m \to T_i^{*,m}$ so that the following holds: For each $t_i \in T_i^m$, $\delta_i^n(t_i) = \delta_i^{*,n}(f_i^{m,n}(t_i))$.*

Before coming to the proof of Lemma B.9, let us note that the Lemma delivers an $f_i^{m,n} : T_i^m \to T_i^{*,m}$ that is Borel measurable and preserves n^{th}-order beliefs. This follows since T_i^m is countable. The fact that $f_{-i}^{m,n}$ is measurable is important in showing the existence of the map $f_i^{m+1,n+1}$.

Proof. The structure of the proof is as follows: We fix a type $t_i \in T_i^m$ and show that there exists a type $t_i^* \in T_i^m$ with $\delta_i^{*,n}(t_i^*) = \delta_i^n(t_i)$. The map $f_i^{m,n} : T_i^m \to T_i^{*,m}$ can then be constructed by setting $f_i^{m,n}(t_i)$ to be the associated t_i^*. The proof is by induction on n.

$\mathbf{n = 1}$: First consider $m = 1$ and let $f_i^{1,1} : T_i^1 \to T_i^{*,1}$ be an arbitrary map. Since $t_i \in T_i^1$ and $f_i^{1,1}(t_i) \in T_i^{*,1}$ are both 1-types in their respective type structures, it follows that

$$\text{marg}_{S_{-i}} \beta_i(t_i) = \mu_i = \text{marg}_{S_{-i}} \beta_i^*(f_i^{1,1}(t_i)).$$

By Lemma B.1, $\delta_i^1(t_i) = \text{marg}_{S_{-i}}\beta_i(t_i)$ and $\delta_i^{*,1}(f_i^{1,1}(t_i)) = \text{marg}_{S_{-i}}\beta_i^*(f_i^{1,1}(t_i))$. From this, the claim follows.

Next consider $m \geq 2$. Fix some $t_i \in T_i^m$. Note, there exists some $\nu_i \in \Delta(S_{-i} \times T_{-i}^*)$ so that $\text{marg}_{S_{-i}}\nu_i = \text{marg}_{S_{-i}}\beta_i(t_i)$ and $\nu_i(S_{-i} \times T_{-i}^{*,m-1}) = 1$. By construction, there exists some $t_i^* \in T_i^{*,m}$ so that $\beta_i^*(t_i^*) = \nu_i$. Now notice that

$$\delta_i^1(t_i) = \text{marg}_{S_{-i}}\beta_i(t_i) = \text{marg}_{S_{-i}}\nu_i = \delta_i^{*,1}(t_i^*).$$

(The first and last equality follows from Lemma B.1. The middle equality comes from the definition of ν_i.) From this, the claim follows.

$\mathbf{n \geq 2}$: Suppose the claim holds for $n \geq 1$. We show that it also holds for $n+1$.

First consider $m = 1$. Note, by the induction hypothesis, for each player j, there exists a mapping $f_j^n : T_j \to T_j^*$ so that $f_j^n(t_j) = f_j^{m,n}(t_j)$ for some m with $t_j \in T_j^m$. (Note, the choice of m does not matter–we only require that $t_j \in T_j^m$.) So the product map $f_{-i}^n : T_{-i} \to T_{-i}^*$ satisfies the following property:

$$\rho_i^{n+1}(s_{-i}, t_{-i}) = \rho_i^{*,n+1}(s_{-i}, f_{-i}^n(t_{-i})).$$

(This uses Lemmata B.2-B.3.) Thus, for each event $E_{-i}^{n+1} \subseteq X_i^n \times H_{-i}^n$,

$$(\rho_i^{n+1})^{-1}(E_{-i}^{n+1}) = (\text{id}_{-i} \times f_{-i}^n)^{-1}((\rho_i^{*,n+1})^{-1}(E_{-i}^{n+1})), \tag{1}$$

where $\text{id}_{-i} : S_{-i} \to S_{-i}$ is the identity map.

Fix some $t_i \in T_i^1$. Let $\nu_i \in \Delta(S_{-i} \times T_{-i}^*)$ be the image measure of $\beta_i(t_i)$ under $(\text{id}_{-i} \times f_{-i}^n)$. By construction, there exists a type $t_i^* \in T_i^{*,1}$ with $\beta_i^*(t_i^*) = \nu_i$. It remains to show that $\delta_i^{*,n+1}(t_i^*) = \delta_i^{n+1}(t_i)$.

Fix some event $E_{-i}^{n+1} \subseteq X_i^n \times H_{-i}^n$. Note,

$$\begin{aligned}
\delta_i^{*,n+1}(t_i^*)(E_{-i}^{n+1}) &= \nu_i\left((\rho_i^{*,n+1})^{-1}(E_{-i}^{n+1})\right) \\
&= \beta_i(t_i)\left((\text{id}_{-i} \times f_{-i}^n)^{-1}\left((\rho_i^{*,n+1})^{-1}(E_{-i}^{n+1})\right)\right) \\
&= \beta_i(t_i)\left((\rho_i^{n+1})^{-1}(E_{-i}^{n+1})\right) \\
&= \delta_i^{n+1}(t_i))(E_{-i}^{n+1}),
\end{aligned}$$

where the third line uses Equation 1. This establishes $\delta_i^{*,n+1}(t_i^*) = \delta_i^{n+1}(t_i)$.

Next consider $m \geq 2$. By the induction hypothesis and Lemmata B.2-B.3, for each $t_{-i} \in T_{-i}^{m-1}$,

$$\rho_i^{n+1}(s_{-i}, t_{-i}) = \rho_i^{*,n+1}(s_{-i}, f_{-i}^{m-1,n}(t_{-i})).$$

Thus, for each event $E_{-i}^{n+1} \subseteq X_i^n \times H_{-i}^n$,

$$(\rho_i^{n+1})^{-1}(E_{-i}^{n+1}) \cap (S_{-i} \times T_{-i}^{m-1}) = (\text{id}_{-i} \times f_{-i}^{m-1,n})^{-1}((\rho_i^{*,n+1})^{-1}(E_{-i}^{n+1}) \cap (S_{-i} \times T_{-i}^{*,m-1})). \tag{2}$$

Fix some $t_i \in T_i^m$. Let $\nu_i \in \Delta(S_{-i} \times T_{-i}^*)$ satisfy the following: For each $E_{-i}^* \subseteq S_{-i} \times T_{-i}^{*,m-1}$,

$$\nu_i(E_{-i}^*) = \beta_i(t_i)((\text{id}_{-i} \times f_{-i}^{m-1,n})^{-1}(E_{-i}^*))$$

and $\nu_i(S_{-i} \times (T^*_{-i} \backslash T^{*,m-1}_{-i})) = 0$. Since $\beta_i(t_i)(S_{-i} \times T^{m-1}_{-i}) = 1$, this is a well-defined measure in $\Delta(S_{-i} \times T_{-i})$ and, moreover, $\nu_i(S_{-i} \times T^{*,m-1}_{-i}) = 1$. By construction, there exists a type $t^*_i \in T^{*,m}_i$ with $\beta^*_i(t^*_i) = \nu_i$. It remains to show that $\delta^{*,n+1}_i(t^*_i) = \delta^{n+1}_i(t_i)$.

Fix some event $E^{n+1}_{-i} \subseteq X^n_i \times H^n_{-i}$. Note,

$$\begin{aligned}
\delta^{*,n+1}_i(t^*_i)(E^{n+1}_{-i}) &= \nu_i((\rho^{*,n+1}_i)^{-1}(E^{n+1}_{-i})) \\
&= \nu_i\left(((\rho^{*,n+1}_i)^{-1}(E^{n+1}_{-i})) \cap (S_{-i} \times T^{*,m-1}_{-i})\right) \\
&= \beta_i(t_i)\left((\mathrm{id}_{-i} \times f^{m-1,n}_{-i})^{-1}\left((\rho^{*,n+1}_i)^{-1}(E^{n+1}_{-i})) \cap (S_{-i} \times T^{*,m-1}_{-i})\right)\right) \\
&= \beta_i(t_i)\left((\rho^{n+1}_i)^{-1}(E^{n+1}_{-i}) \cap (S_{-i} \times T^{m-1}_{-i})\right) \\
&= \beta_i(t_i)\left((\rho^{n+1}_i)^{-1}(E^{n+1}_{-i})\right) \\
&= \delta^{n+1}_i(t_i)(E^{n+1}_{-i}),
\end{aligned}$$

where the second line uses the fact that $\nu_i(S_{-i} \times T^{*,m-1}_{-i}) = 1$, the third line follows from the construction of ν_i, the fourth line follows from Equation 2, and the fifth line uses the fact that $\beta_i(t_i)(S_{-i} \times T^{m-1}_{-i}) = 1$. This establishes $\delta^{*,n+1}_i(t^*_i) = \delta^{n+1}_i(t_i)$. ■

Remark B.1. [**TODO:** Insert discussion of the difficulty in extending the previous result to any level-k type structure]

Appendix C Proofs for Sections 5-6

C.1 Proofs for Section 5

Lemma C.1. *Let $J_i : S_i \twoheadrightarrow \Delta(S_{-i})$ be a correspondence with*

$$J_i(s_i) = \{\nu_i \in \Delta(S_{-i}) : s_i \in \mathbb{BR}_i[\nu_i]\}.$$

Then $J_i(s_i)$ is closed-valued. Moreover, if $s_i \in S^1_i$, then $J_i(s_i)$ is non-empty valued.

Proof. Let $\hat{\pi}_i : S_i \times \Delta(S_{-i}) \to \mathbb{R}$ be defined by

$$\hat{\pi}_i(s_i, \nu_i) = \sum_{S_{-i}} \pi_i(s_i, s_{-i})\nu_i(s_{-i}).$$

It follows from Theorem 15.3 in Aliprantis and Border (2007) and the fact that S_i is finite that $\hat{\pi}_i$ is continuous. Moreover, since $S_i \times \Delta(S_{-i})$ is compact, $\hat{\pi}_i$ is bounded. As a consequence, the function $\tilde{\pi}_i : S_i \times S_i \times \Delta(S_{-i}) \to \mathbb{R}$ defined by

$$\tilde{\pi}_i(s_i, r_i, \nu_i) = \hat{\pi}_i(s_i, \nu_i) - \hat{\pi}_i(r_i, \nu_i)$$

is continuous and bounded.

Now, fix a sequence $(\nu^1_i, \nu^2_i, \ldots)$ with each $\nu^k_i \in J_i(s_i)$. Then, for each ν^k_i and each $r_i \in S_i$, $\tilde{\pi}_i(s_i, r_i, \nu^k_i) \geq 0$. If $(\nu^1_i, \nu^2_i, \ldots)$ converges to ν_i then, for each $r_i \in S_i$, $\tilde{\pi}_i(s_i, r_i, \nu_i) \geq 0$. (See Theorem 15.3 in Aliprantis and Border, 2007, which uses the fact that $\tilde{\pi}_i$ is continuous and bounded.) Thus, $\nu_i \in J_i(s_i)$ and $J_i(s_i)$ is closed. ■

Lemma C.2.

(i) If E_{-i} is Borel then $B_i(E_{-i})$ is Borel.

(ii) If $E_{-i} = \emptyset$, then $B_i(E_{-i}) = \emptyset$ and so Borel.

Proof. Part (i) follows from Lemma 15.16 in Aliprantis and Border (2007) and the fact that β_i is measurable. Part (ii) is immediate. ■

Lemma C.3. *For each m, the sets R_i^m are Borel.*

Proof. The proof is by induction on m.

$m = 1$: Fix a strategy s_i and let

$$O(s_i) = \{\nu_i \in \Delta(S_{-i} \times T_{-i}) : s_i \in \mathbb{BR}_i[\operatorname{marg}_{S_{-i}} \nu_i]\}.$$

By Lemma C.1 and Lemma A.1, $O[s_i]$ is closed. From this and the fact that β_i is measurable, each $\{s_i\} \times \beta_i^{-1}(\hat{O}[s_i])$ is Borel. Now observe that

$$R_i^1 = \bigcup_{s_i \in S_i} \left(\{s_i\} \times \beta_i^{-1}(O(s_i))\right)$$

and, therefore, R_i^1 is Borel.

$m \geq 2$: Assume that, for each i, R_i^m is Borel. As such, each R_{-i}^m is also Borel. So by Lemma C.2(i), R_i^m is Borel. ■

C.2 Proof of Theorem 6.1

Proposition C.1. *Fix an epistemic game $(G, \mathcal{T})$ where $\mathcal{T}$ is a level-k type structure for μ. Then:*

(i) $\operatorname{proj}_{S_i}\left(R_i^1 \cap (S_i \times T_i^1)\right) = L_i^1[\mu]$, and

(ii) for each $k \geq 1$, $\operatorname{proj}_{S_i}\left(R_i^k \cap (S_i \times T_i^m)\right) \subseteq L_i^m$.

Proof. Begin with part (i). Fix some $s_i \in \operatorname{proj}_{S_i}\left(R_i^1 \cap (S_i \times T_i^1)\right)$. Then there exists some $t_i \in T_i^1$ so that $(s_i, t_i) \in R_i^1$. As such, $s_i \in \mathbb{BR}_i[\operatorname{marg}_{S_{-i}}\beta_i(t_i)]$ and $\operatorname{marg}_{S_{-i}}\beta_i(t_i) = \mu_i$. So $s_i \in L_i^1[\mu]$. Conversely, fix $s_i \in L_i^1[\mu]$. Then $s_i \in \mathbb{BR}_i[\mu_i]$ and, for each $t_i \in T_i^1$, $\operatorname{marg}_{S_{-i}}\beta_i(t_i) = \mu_i$. Thus, $\{s_i\} \times T_i^1 \subseteq R_i^1 \cap (S_i \times T_i^1)$. As such, $L_i^1[\mu] \subseteq \operatorname{proj}_{S_i}\left(R_i^1 \cap (S_i \times T_i^1)\right)$.

The proof of part (ii) is by induction on m. The case of $m = 1$ follows from part (i). Assume the claim holds for m. Fix some $s_i \in \operatorname{proj}_{S_i}\left(R_i^{m+1} \cap (S_i \times T_i^{m+1})\right)$. Then there exists some $t_i \in T_i^{m+1}$ so that $(s_i, t_i) \in R_i^{m+1}$. As such, $s_i \in \mathbb{BR}_i[\operatorname{marg}_{S_{-i}}\beta_i(t_i)]$. Moreover, $\beta_i(t_i)(R_{-i}^m \cap (S_{-i} \times T_{-i}^m)) = 1$. So, by the induction hypothesis, $\operatorname{marg}_{S_{-i}}\beta_i(t_i)(L_{-i}^m) = 1$. As such, $s_i \in L_i^{m+1}$. ■

Proof of Theorem 6.1. Part (i) is Proposition C.1. So we focus on part (ii). Throughout, fix a complete level-k type structure for $\boldsymbol{\mu}$ with covers $\mathcal{C}_i = \{T_i^m : m = 1, 2, \ldots\}$ satisfying conditions (i)-(ii)-(iii) of Definition 4.2. The proof is by induction on m.

The case of $m = 1$ is part (i) of Proposition C.1. So, assume the result holds for m. By part (ii) of Proposition C.1, it suffices to show that

$$L_i^{m+1} \subseteq \text{proj}_{S_i}\left(R_i^{m+1} \cap (S_i \times T_i^{m+1})\right).$$

Fix $s_i \in L_i^{m+1}$. Then there exists some $\nu_i \in \Delta(S_{-i})$ such that $s_i \in \mathbb{BR}_i[\nu_i]$, and $\nu_i(L_{-i}^m) = 1$. We will use ν_i to construct a $\hat{\nu}_i \in \Delta(S_{-i} \times T_{-i})$ so that: (i) $\text{marg}_{S_{-i}}\hat{\nu}_i = \nu_i$, (ii) $\hat{\nu}_i(S_{-i} \times T_{-i}^m) = 1$, and (iii) for each $n \leq m$, $\hat{\nu}_i(R_{-i}^n) = 1$. We then show that this suffices to deliver the result.

Step 1: By the induction hypothesis, for each player j, there exists a mapping $\tau_j^m : L_j^m \to T_j^m$ that satisfies the following property: For each $s_j \in L_j^m$, $(s_j, \tau_j^m(s_j)) \in R_j^m \cap (S_j \times T_j^m)$. Let $\tau_{-i}^m : L_{-i}^m \to T_{-i}^m$ be the associated product map. For each $s_{-i} \in L_{-i}^m$, set $\hat{\nu}(s_{-i}, \tau_{-i}^m(s_{-i})) = \nu(s_{-i})$ and, for each $(s_{-i}, t_{-i}) \in S_{-i} \times T_{-i} \backslash (\text{gr}(\tau_{-i}^m))$, set $\hat{\nu}(s_{-i}, t_{-i}) = 0$. This gives a $\hat{\nu}_i \in \Delta(S_{-i} \times T_{-i})$. By the construction and the fact that T_{-i}^m is Borel, we have $\hat{\nu}_i(S_{-i} \times T_{-i}^m) = 1$. By the construction and the fact that each R_{-i}^n is Borel, we have that, for each $n \leq m$, $\hat{\nu}_i(R_{-i}^n) = 1$.

Step 2: By completeness, there exists a type $t_i \in T_{-i}^{m+1}$ with $\beta_i(t_i) = \hat{\nu}_i$. Since $\text{marg}_{S_{-i}}\beta_i(t_i) = \nu_i$ and $s_i \in \mathbb{BR}_i[\nu_i]$, it follows that $(s_i, t_i) \in R_i^1$. Since, for each $n \leq m$, $\beta_i(t_i)(R_{-i}^n) = 1$, $(s_i, t_i) \in R_i^{m+1}$. ■

C.3 Result for Section 6.3

Lemma C.4. *Fix an anchor $\boldsymbol{\mu}$. For each $m \geq 1$ and each $n \geq m$, $L_i^n \subseteq S_i^m$.*

Proof. The proof is by induction on m. For $m = 1$ and each $n \geq 1$, it is immediate that $L_i^n \subseteq S_i^1$; thus, the result holds for $m = 1$. Suppose the result holds for $m \geq 1$. Fix $n \geq m$ and note that $s_i \in L_i^{n+1}$ if and only if s_i is a best response under some $\nu_i \in \Delta(S_{-i})$ with $\nu_i(L_{-i}^n) = 1$. By the induction hypothesis, $L_{-i}^n \subseteq S_{-i}^m$ and so $\nu_i(S_{-i}^m) = 1$. Thus, $s_i \in S_{-i}^{m+1}$. ■

Appendix D Proofs for Section 8

D.1 Proof of Proposition 8.1

The proof is analogous to Example 8.1: Since each $|S_i| \geq 2$, take $\{\Box_i, \Diamond_i\} \subseteq S_i$. Fix a non-degenerate anchor $\boldsymbol{\mu}$, i.e., an anchor where each μ_i does not assign probability 1 to some strategy. Then, for each i, there exists some strategy $s_{-i} \in S_{-i}$ so that $\mu_i(s_{-i}) \in (0, 1)$. Without loss of generality, suppose that, for each i, this strategy profile is $\Box_{-i}$.

Inductively define $h_{i,\Box}^m$ so that $h_{i,\Box}^1(\Box_{-i}) = 1$ and $h_{i,\Box}^{m+1}(\Box_{-i}, \ldots, h_{-i,\Box}^m) = 1$. Set $h_{i,\Box} = (h_{i,\Box}^1, h_{i,\Box}^2, \ldots)$. Likewise, for each player i, inductively define h_i^m as follows: First, set $h_i^1 = \mu_i$. Second, $h_i^m(\Box_{-i}, h_{-i,\Box}^1, \ldots, h_{-i,\Box}^m) = p \in (0, \mu_i(\Box_{-i})]$. (Note, p does not depend on m.) Set $h_i = (h_i^1, h_i^2, \ldots)$. Proposition 8.1 will follow from the following two Lemmata.

Lemma D.1. *Fix a type structure $\mathcal{T} = (S_{-i}, T_i, \beta_i : i \in I)$. If there exists a type $t_i \in T_i$ with $\delta_i(t_i) = h_i$, then there must be a type $t_{-i,\Box} \in T_{-i}$ with $\delta_{-i}(t_{-i,\Box}) = h_{-i,\Box}$.*

Proof. Suppose there is a type $t_i \in T_i$ with $\delta_i(t_i) = h_i$. Note, for each $m \geq 1$,

$$h_i^{m+1}(\square_{-i}, h^1_{-i,\square}, \ldots, h^m_{-i,\square}) = p$$

or, equivalently, $\beta_i(t_i)(E_i^{m+1}) = p$ for

$$E_i^{m+1} := (\rho_i^{m+1})^{-1}(\{\square_{-i}, h^1_{-i,\square}, \ldots, h^m_{-i,\square}\}).$$

Observe that the sets E_i^m are decreasing, i.e., for each $m \geq 2$, $E_i^{m+1} \subseteq E_i^m$. Since $(\beta_i(t_i)(E_i^m) : m \geq 2) = (p, p, p, \ldots)$,

$$p = \lim_{m\to\infty} \beta_i(t_i)(E_i^m) = \beta_i(t_i)(\bigcap_{m\geq 2} E_i^m).$$

(See, e.g., Theorem 10.8 in Aliprantis and Border (2007).) Thus,

$$\bigcap_{m\geq 2} E_i^m \neq \emptyset,$$

i.e., there exists some type $t_{-i} \in T_{-i}$ with $\delta_{-i}(t_{-i}) = (h^1_{-i,\square}, h^2_{-i,\square}, \ldots)$, as required. ■

Lemma D.2. *If $\mathcal{T} = (S_{-i}, T_i, \beta_i : i \in I)$ is a level-k type structure for $\boldsymbol{\mu}$, then there is no type $t_{i,\square} \in T_i$ with $\delta_i(t_{i,\square}) = h_{i,\square}$.*

Proof. For each $i \in I$, let $\mathcal{C}_i = \{T_i^m : m = 1, 2, \ldots\}$ be a Borel cover so that $(\mathcal{C}_i : i \in I)$ jointly satisfy conditions (i)-(ii) of Definition 4.1. We will show that, for each $m \geq 1$, and each $t_i \in T_i^m$, $\delta_i^m(t_i) \neq h^m_{i,\square}$. The proof is by induction on m.

The case of $m = 1$ is immediate: If $t_i \in T_i^1$, $\delta_i^1(t_i)(\square_{-i}) \neq 1$ and so $\delta_i^1(t_i) \neq h^1_{i,\square}$. Suppose then that the claim holds for m. Fix $t_i \in T_i^{m+1}$. By the induction hypothesis,

$$(\rho_i^{m+1})^{-1}(S_{-i} \times \{(h^1_{-i,\square}, \ldots, h^m_{-i,\square})\}) \cap (S_{-i} \times T^m_{-i}) = \emptyset.$$

Since $\beta_i(t_i)(S_{-i} \times T^m_{-i}) = 1$,

$$\beta_i(t_i)((\rho_i^{m+1})^{-1}(S_{-i} \times \{(h^1_{-i,\square}, \ldots, h^m_{-i,\square})\})) = 0$$

and so $\delta_i^{m+1}(t_i) \neq h^{m+1}_{i,\square}$. ■

D.2 Properties of Level-k Type Structures

Example D.1. This example shows that, for a given level-k type structure, we may not be able to choose the cover to be a partition. As such, we may have that a type is both a k-type and an ℓ-type for every associated cover.

Construct an S-based level-k type structures for μ, viz. $\mathcal{T} = (S_{-i}, T_i, \beta_i : i \in I)$, as follows: For each i, take $T_i = \mathbb{N}_+$. Choose β_i so that it satisfies the following properties. First, $\operatorname{marg}_{S_{-i}} \beta_i(t_i) = \mu_i$ if and only if $t_i \in \{1, 3\}$. Second, $\operatorname{Supp} \operatorname{marg}_{T_{-i}} \beta_i(1) = T_{-i}$. Third, $\beta_i(2)(S_{-i} \times \{1\}) = \beta_i(2)(S_{-i} \times \{3\}) = \frac{1}{2}$. Fourth, for each $k \geq 2$, $\beta_i(k+1)(S_{-i} \times \{k\}) = 1$.

This is a level-k type structure for μ. We can choose the cover $\{T_i^m : k = 1, 2, \ldots\}$ so that $T_i^1 = \{1, 3\}$ and, for each $k \geq 2$, $T_i^m = \{k\}$. This cover is non-partitional. However, any cover must be non-partitional. To see this, fix a cover $\{U_i^k : k = 1, 2, \ldots\}$. Since $\operatorname{Supp} \operatorname{marg}_{T_{-i}} \beta_i(1) = T_{-i}$, it must be that $1 \in U_i^1$. So, U_i^1 is either $\{1\}$ or $\{1, 3\}$. If $U_i^1 = \{1\}$ then $U_{-i}^2 = \emptyset$. So we must have $U_i^1 = \{1, 3\}$ and, from this, it follows that $U_i^2 = \{2\}$. But this implies that $U_i^3 = \{3\}$. Thus, any cover must have $U_i^1 \cap U_i^3 \neq \emptyset$. $\square$

Example D.2. This example shows that, for any anchor $\boldsymbol{\mu}$, there may be a level-k type structure for $\boldsymbol{\mu}$ where the associated Borel cover is not unique. As a result, a type t_i may be a k-type for one associated cover and an ℓ-type for another associated cover, despite the fact that $k \neq \ell$.

Fix an anchor $\boldsymbol{\mu}$. Construct a type structure as follows: For each i, take $T_i = \mathbb{N}_+$. Choose β_i so that it satisfies the following properties. First, $\operatorname{marg}_{S_{-i}} \beta_i(t_i) = \mu_i$ if and only if $t_i \in \{1, 3\}$. Second, $\operatorname{Supp} \operatorname{marg}_{T_{-i}} \beta_i(1) = T_{-i}$. Third, for each $k \geq 1$, $\beta_i(k+1)(S_{-i} \times \{k\}) = 1$.

This is a level-k type structure for $\boldsymbol{\mu}$. Notice, we can take the cover $\{T_i^m : k = 1, 2, \ldots\}$ so that $T_i^m = \{k\}$ for each k. This cover is a partition. However, there is a second non-partitional cover $\{U_i^k : k = 1, 2, \ldots\}$ with $U_i^1 = \{1, 3\}$ and, for each $k \geq 2$, $U_i^k = \{k\}$. Under the first cover, 3 is a 3-type, while under the second cover, 3 is both a 1-type and a 3-type. $\square$

D.3 Finite-Order Belief Type Structures

Definition D.1. A **finitary S-based type structure** is some $\tilde{\mathcal{T}} = (S_{-i}, \tilde{T}_i, \tilde{\beta}_i : i \in I)$ where,

(i) for each i, $\tilde{T}_i$ is a metrizable set of types for i with $\tilde{T}_i \cap \{d\} = \emptyset$ and

(ii) for each i, $\tilde{\beta}_i : \tilde{T}_i \to \Delta(S_{-i} \times \tilde{T}_{-i}) \cup \{d\}$ is a measurable belief map for i.

Say $(s_i, \tilde{t}_i)$ is **rational** if $\tilde{\beta}_i(\tilde{t}_i) \in \Delta(S_{-i} \times (\tilde{T}_{-i} \cup \{d\}))$ and satisfies the condition in Definition 5.1. Say $\tilde{t}_i$ **believes** an event E_{-i} if $\tilde{\beta}_i(\tilde{t}_i) \in \Delta(S_{-i} \times (\tilde{T}_{-i} \cup \{d\}))$ and $\tilde{t}_i$ satisfies the condition in Definition 5.2.

We define RmBR analogously to Definition 5.3. In particular, we write $\tilde{R}_i^1$ for the set of rational strategy-type pairs and $\tilde{R}_i^{m+1}$ for the set of strategy-type pairs which satisfy rationality and m^{th}-order belief of rationality.

Each ordinary type structure is also a finitary S-based type structure. With this in mind, we focus on showing that the RmBR predictions of a finitary type structure can be replicated in an ordinary type structure. In doing so, we will focus on type structures that are first-order complete: Call $\tilde{\mathcal{T}}$ **first-order complete** if, for each $\nu_i \in \Delta(S_{-i})$, there exists some $\tilde{t}_i \in \tilde{T}_i$ with $\operatorname{marg}_{S_{-i}} \tilde{\beta}_i(\tilde{t}_i) = \nu_i$.

Proposition D.1. *Fix a game with no weakly dominant strategy. Let $\tilde{\mathcal{T}} = (S_{-i}, \tilde{T}_i, \tilde{\beta}_i : i \in I)$ be a finitary S-based type structure that is first-order complete. Then, there exists an ordinary S-based type structure $\mathcal{T} = (S_{-i}, T_i, \beta_i : i \in I)$ with each $T_i \subseteq \tilde{T}_i$ so that*

(i) for each $t_i \in T_i$, $(s_i, t_i) \in R_i^m$ if and only if $(s_i, t_i) \in \tilde{R}_i^m$, and

(ii) $\operatorname{proj}_{S_i} R_i^m = \operatorname{proj}_{S_i} \tilde{R}_i^m$.

To prove Proposition D.1, we will make use of the following fact: If a game has no weakly dominant strategy for i, then we can find a mapping $f_i : S_i \to \Delta(S_{-i})$ so that, for each $s_i \in S_i$, $s_i \notin \mathbb{BR}_i[f_i(s_i)]$. We make use of these mappings below.

Proof of Proposition D.1. Fix a game with no weakly dominant strategy and an associated finitary S-based type structure that is first-order complete, viz. $\tilde{\mathcal{T}}$. Since there are no weakly dominant strategies, we can find mappings $f_i : S_i \to \Delta(S_{-i})$ so that, for each $s_i \in S_i$, $s_i \notin \mathbb{BR}[f_i(s_i)]$. Since $\mathcal{T}$ is first-order complete, there are mappings $\tau_i : S_i \to \tilde{T}_i$ such that $\text{marg}_{S_{-i}}\tilde{\beta}_i(\tau_i(s_i)) = f_i(s_i)$. Since S_i is finite, τ_i is measurable.

With this background, we can construct $\mathcal{T}$. Let $T_i = \tilde{T}_i \backslash \{t_i \in \tilde{T}_i : \tilde{\beta}_i(t_i) = d\}$. Observe that T_i is a Borel subset of $\tilde{T}_i$. (This follows from the fact that $\tilde{\beta}_i$ is measurable.) Endow T_i with the relative topology and note that it is metrizable.

Observe that, by construction, $\tau_i(S_i) \subseteq T_i$. As such, write $\overline{\tau}_i : S_i \to T_i$ for the restriction of τ_i to the range T_i. Note that $\overline{\tau}_i$ is also measurable. Write $(\text{id}_{-i} \times \overline{\tau}_{-i}) : S_{-i} \times S_{-i} \to S_{-i} \times T_{-i}$ for the associated product mappings. That is, $(\text{id}_{-i} \times \overline{\tau}_{-i})$ is a mapping where, for each $s_{-i} \in S_{-i}$, $(\text{id}_{-i} \times \overline{\tau}_{-i})(s_{-i}, s_{-i}) = (s_{-i}, \overline{\tau}_{-i}(s_{-i}))$. Observe that, since id_{-i} and $\overline{\tau}_{-i}$ are both measurable, $(\text{id}_{-i} \times \overline{\tau}_{-i})$ is measurable.

We now construct β_i. To do so, it will be convenient to derive the mapping from two auxiliary mappings, β_i° and $\beta_i^{\diamond}$. Let T_i° be the set of $t_i \in T_i$ with $\tilde{\beta}_i(t_i)(S_{-i} \times T_{-i}) = 1$. Let $T_i^{\diamond} = T_i \backslash T_i^{\circ}$. Since $\tilde{\beta}_i$ is measurable, both T_i° and $T_i^{\diamond}$ are measurable. Take $\beta_i^{\circ} : T_i^{\circ} \to \Delta(S_{-i} \times T_{-i})$ so that, for each $t_i \in T_i^{\circ}$, $\beta_i^{\circ}(t_i)$ is the restriction of $\tilde{\beta}_i(t_i)$ to $S_{-i} \times T_{-i}$. Note that β_i° is measurable. Take $\beta_i^{\diamond} : T_i^{\diamond} \to \Delta(S_{-i} \times T_{-i})$ so that, for each $t_i \in T_i^{\diamond}$, $\beta_i^{\diamond}(t_i)$ is the image measure of $\text{marg}_{S_{-i}}\tilde{\beta}_i(t_i)$ under $\text{id}_{-i} \times \overline{\tau}_{-i}$. Note, $\beta_i^{\diamond}$ is measurable. Finally, let

$$\beta_i(t_i) = \begin{cases} \beta_i^{\circ}(t_i) & \text{if } t_i \in T_i^{\circ}, \\ \beta_i^{\diamond}(t_i) & \text{if } t_i \in T_i^{\diamond}. \end{cases}$$

Note that β_i is measurable since T_i°, $T_i^{\diamond}$, β_i°, and $\beta_i^{\diamond}$ are each measurable.

Finally, we show that, for each $m \geq 1$ and each $t_i \in T_i$, $(s_i, t_i) \in R_i^m$ if and only if $(s_i, t_i) \in \tilde{R}_i^m$. This will imply that, for each $m \geq 1$, $\text{proj}_{S_i} R_i^m = \text{proj}_{S_i}(\tilde{R}_i^m \cap (S_i \times T_i))$. Now observe that, for each $m \geq 1$, $\text{proj}_{S_i}(\tilde{R}_i^m \cap (S_i \times T_i)) = \text{proj}_{S_i}\tilde{R}_i^m$. As such, for each $m \geq 1$, $\text{proj}_{S_i} R_i^m = \text{proj}_{S_i}\tilde{R}_i^m$.

In fact, we will show a slightly stronger claim:

(i) For each $m \geq 1$ and each $t_i \in T_i$, $(s_i, t_i) \in R_i^m$ if and only if $(s_i, t_i) \in \tilde{R}_i^m$.

(ii) For each $m \geq 2$ and each $t_i \in T_i^{\diamond}$, $S_i \times \{t_i\} \cap R_i^m = \emptyset$ and $S_i \times \{t_i\} \cap \tilde{R}_i^m = \emptyset$.

The proof is by induction on m.

$m = 1$: Fix $t_i \in T_i$. By construction, $\text{marg}_{S_{-i}}\tilde{\beta}_i(t_i) = \text{marg}_{S_{-i}}\beta_i(t_i)$. As such, $(s_i, t_i) \in R_i^1$ if and only if $(s_i, t_i) \in \tilde{R}_i^1$.

$m = 2$: Fix $t_i \in T_i$. If $t_i \in T_i^{\circ}$, then t_i believes R_{-i}^1 if and only if t_i believes $\tilde{R}_{-i}^1$. (This follows from the construction.) If $t_i \in T_i^{\diamond}$, then t_i does not believe $\tilde{R}_{-i}^1$. (This follows from the fact that $\tilde{R}_{-i}^1 \cap (S_{-i} \times \tilde{T}_{-i} \backslash T_{-i}) = \emptyset$.) Thus, we must show that t_i does not believe R_{-i}^1. To see this, observe that $\beta_i(t_i)(S_{-i} \times \tau_{-i}(S_{-i})) = 1$ and, by construction, $(S_{-i} \times \tau_{-i}(S_{-i})) \cap R_{-i}^1 = \emptyset$. As such, t_i does not believe R_{-i}^1.

$m \geq 3$: Assume the claim holds for $m \geq 3$ and we show that it also holds for $m + 1$. Fix $t_i \in T_i$. If $t_i \in T_i^{\circ}$, then t_i believes R_{-i}^m if and only if t_i believes $\tilde{R}_{-i}^m$. (This follows from the construction.) If $t_i \in T_i^{\diamond}$, then by the induction hypothesis, $(S_i \times \{t_i\}) \cap R_i^m = \emptyset$ and $(S_i \times \{t_i\}) \cap \tilde{R}_i^m = \emptyset$. As such, if $t_i \in T_i^\diamond$, $(S_i \times \{t_i\}) \cap R_i^{m+1} = \emptyset$ and $(S_i \times \{t_i\}) \cap \tilde{R}_i^{m+1} = \emptyset$. ■

D.4 RmBR Behavior of k-Types

Lemma D.3. *Fix an epistemic game $(G, \mathcal{T})$, where $\mathcal{T}$ that is a level-k type structure for μ. If, for each i, $\mu_i(S_{-i} \backslash S_{-i}^1) > 0$, then*

$$\bigcup_{k \geq m} \left(R_i^m \cap (S_i \times T_i^k)\right) = \bigcup_{k \geq 1} \left(R_i^m \cap (S_i \times T_i^k)\right)$$

for each m.

Proof. The proof is by induction on m. For $m = 1$, the claim is immediate. So suppose that $m \geq 2$. We will show that, for each $k < m$ and each $(s_i, t_i) \in S_i \times T_i^k$, $(s_i, t_i) \notin R_i^{k+1}$. From this it follows that $(s_i, t_i) \notin R_i^m$ and so $R_i^m \cap (S_i \times T_i^k) = \emptyset$.

The proof is by induction on k. For $(s_i, t_i) \in S_i \times T_i^1$, $\operatorname{marg}_{S_{-i}} \beta_i(t_i)(S_{-i} \backslash S_{-i}^1) > 0$ and so $(s_i, t_i) \notin R_i^2$. Assume that the claim holds for $k \leq m - 2$. If $(s_i, t_i) \in S_i \times T_i^{k+1}$, $\beta_i(t_i)(S_{-i} \times T_{-i}^k) = 1$ and so, by the induction hypothesis, $\beta_i(t_i)(R_{-i}^{k+1}) = 0$. Thus, $(s_i, t_i) \notin R_i^{k+2}$. ■

Proof of Lemma 8.1. By Lemma D.3,

$$\bigcup_{k \geq m} \left(R_i^m \cap (S_i \times T_i^k)\right) = \bigcup_{k \geq 1} \left(R_i^m \cap (S_i \times T_i^m)\right) = R_i^m,$$

from which the claim follows. ■

References

Alaoui, Larbi and Antonio Penta. 2016. “Endogenous depth of reasoning.” *The Review of Economic Studies* 83(4):1297–1333.

Alaoui, Larbi, Katharina A Janezic and Antonio Penta. 2020. “Reasoning about others’ reasoning.” *Journal of Economic Theory* 189:105091.

Aliprantis, C.D. and K.C. Border. 2007. *Infinite dimensional analysis: a hitchhiker’s guide.* Springer Verlag.

Battigalli, P., A. Friedenberg and M. Siniscalchi. 2012. *Strategic Uncertainty: An Epistemic Approach to Game Theory.* (Working Title).

Battigalli, P. and M. Siniscalchi. 2002. “Strong Belief and Forward Induction Reasoning.” *Journal of Economic Theory* 106(2):356–391.

Brandenburger, A. and E. Dekel. 1987. “Rationalizability and Correlated Equilibria.” *Econometrica* 55(6):1391–1402.

Brandenburger, A. and E. Dekel. 1993. “Hierarchies of Beliefs and Common Knowledge.” *Journal of Economic Theory* 59:189–189.

Brandenburger, Adam, Alexander Danieli and Amanda Friedenberg. 2021. "The implications of finite-order reasoning." *Theoretical Economics* 16(4):1605–1654.

Brocas, Isabelle, Juan D Carrillo, Stephanie W Wang and Colin F Camerer. 2014. "Imperfect choice or imperfect attention? Understanding strategic thinking in private information games." *Review of Economic Studies* 81(3):944–970.

Camerer, C., T. Ho and J. Chong. 2004. "A Cognitive Hierarchy Model of Games." *The Quarterly Journal of Economics* 119(3):861–898.

Cooper, David J, Enrique Fatas, Antonio J Morales and Shi Qi. 2024. "Consistent depth of reasoning in level-k models." *American Economic Journal: Microeconomics* 16(4):40–76.

Costa-Gomes, M., V. Crawford and B. Broseta. 2001. "Cognition and Behavior in Normal-Form Games: An Experimental Study." *Econometrica* 69(5):1193–1235.

Costa-Gomes, Miguel A and Vincent P Crawford. 2006. "Cognition and behavior in two-person guessing games: An experimental study." *American economic review* 96(5):1737–1768.

Crawford, Vincent P and Nagore Iriberri. 2007. "Fatal attraction: Salience, naivete, and sophistication in experimental "hide-and-seek" games." *American Economic Review* 97(5):1731–1750.

Friedenberg, A. 2010. "When do Type Structures Contain all Hierarchies of Beliefs?" *Games and Economic Behavior* 68(1):108–129.

Friedenberg, A. and T. Kneeland. 2024. "Beyond Reasoning about Rationality: Evidence of Strategic Reasoning.". https://www.amandafriedenberg.org/working-papers.

Friedenberg, Amanda and H Jerome Keisler. 2021. "Iterated dominance revisited." *Economic Theory* 72(2):377–421.

Georganas, Sotiris, Paul J Healy and Roberto A Weber. 2015. "On the persistence of strategic sophistication." *Journal of Economic Theory* 159:369–400.

Ghosh, Sujata, Aviad Heifetz and Rineke Verbrugge. 2016. "Do players reason by forward induction in dynamic perfect information games?" *arXiv preprint arXiv:1606.07521* .

Ghosh, Sujata and Rineke Verbrugge. 2018. "Studying strategies and types of players: Experiments, logics and cognitive models." *Synthese* 195(10):4265–4307.

Harsanyi, J.C. 1967. "Games with Incomplete Information Played by "Bayesian" Players, I-III. Part I. The Basic model." *Management Science* pp. 159–182.

Healy, Paul J. 2024. "Epistemic experiments: Utilities, beliefs, and irrational play." *Unpublished manuscript, Ohio State University, Columbus, OH* .

Heifetz, A. and D. Samet. 1998. "Topology-free typology of beliefs." *Journal of Economic Theory* 82(2):324–341.

Heifetz, A. and W. Kets. 2018. "Robust Multiplicity with a Grain of Naivite." *Theoretical Economics* 13:415–465.

Kets, W. 2010. "Bounded Reasoning and Higher-Order Uncertainty." http://tuvalu.santafe.edu/~willemien.kets/.

Kneeland, Terri. 2015. "Identifying Higher-Order Rationality." *Econometrica* 83:2065–2079.

Li, Ying Xue and Burkhard C Schipper. 2020. "Strategic reasoning in persuasion games: An experiment." *Games and Economic Behavior* 121:329–367.

Liu, Shuige and Gabriel Ziegler. 2025. "Reasoning about Bounded Reasoning." *arXiv preprint arXiv:2506.19737* .

Mertens, J.F. and S. Zamir. 1985. "Formulation of Bayesian Analysis for Games with Incomplete Information." *International Journal of Game Theory* 14(1):1–29.

Nagel, R. 1995. "Unraveling in Guessing Games: An Experimental Study." *The American Economic Review* 85(5):1313–1326.

Schipper, Burkhard C and Hang Zhou. 2024. "Level-k thinking in the extensive form." *Economic Theory* pp. 1–41.

Stahl, D. and P. Wilson. 1995. "On Players' Models of Other Players: Theory and Experimental Evidence." *Games and Economic Behavior* 10(1):218–254.

Stahl, Dale O. and Paul W. Wilson. 1994. "Experimental Evidence on Player's Models of Other Players." *Journal of Economic Behavior and Organization* 25(3):309–327.

Strzalecki, Tomasz. 2014. "Depth of reasoning and higher order beliefs." *Journal of Economic Behavior & Organization* 108:108–122.

Tan, T.C.C. and S.R. Werlang. 1988. "The Bayesian Foundations of Solution Concepts of Games." *Journal of Economic Theory* 45(2):370–391.

Wright, James R and Kevin Leyton-Brown. 2019. "Level-0 models for predicting human behavior in games." *Journal of Artificial Intelligence Research* 64:357–383.