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

Adam Brandenburger† Amanda Friedenberg‡ Terri Kneeland§

December 22, 2025

Abstract

The level-k concept is widely used to assess players' reasoning in games. This paper argues that the concept can overstate evidence of bounded reasoning. It uses epistemic game theory to model the reasoning process typically associated with level-k behavior. The main theorem shows that the level-k model of reasoning has the predictive power of m-rationalizability. So, behavior viewed as reflecting k levels of reasoning may be consistent with higher levels of reasoning. The paper goes on to provide an epistemic characterization of level-k behavior, which highlights the difficulty in inferring levels of reasoning from the level k categorization.

1 Introduction

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 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

*An earlier version was circulated under the title "Two Approaches to Iterated Reasoning in Games" (December 2020). We thank Andrew Ellis, Kristof Madarasz, Karl Schlag, and 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

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, i.e., 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 $m \geq k$. Thus, a classification of k (according to the level-k model) 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 and the example in Section 2.[1]) 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 2 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 (at least) $(m+1)$ 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

[1] This is also the idea behind Alaoui and Penta's (2016) modification of the 11-20 game.

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 4.2 and 9.A.)

Level-k type structures are engineered to mimic the logic of the level-k model. Despite this, 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. Thus, the reasoning (or "cognitive") process associated with the level-k solution concept has the predictive power of rationalizability. This 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 (i.e., the joint distribution) 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 7.1 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. We discuss this further in Section 7.3.

It is worth emphasizing that the focus of our paper is on the interpretation of the level-k model as a model of reasoning. Our central result shows that the reasoning (or "cognitive") model typically associated with the level-k solution concept has less predictive power than the level-k concept itself. As a consequence, using the concept to evaluate reasoning in games can result in overestimating the extent to which there is bounded reasoning. Of course, this does not question the usefulness of the level-k solution concept for fitting behavior.

Literature This is not the first paper to point to difficulties in drawing inferences about how players reason from the level k classification. 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), providing auxiliary evidence on the anchor (Costa-Gomes and Crawford, 2006; Brocas, Carrillo, Wang and Camerer, 2014; Burchardi and Penczynski, 2014), or designing the game with the aim of making an anchor highly salient (Arad and Rubinstein, 2012). 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 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, Fatas, Morales and Qi, 2024.) Fourth, it may be that the inferred levels of 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 study an idealized setting. In particular, we focus 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, we ignore concerns about portability. We argue that, even in this idealized setting, a classification of level k 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 model a level-0 player as one that has different payoffs from those specified in the game; it then uses rationalizability concepts to analyze the game of incomplete information and to draw connections to the level-k literature. The focus of these papers differs from ours; in particular, they do not directly address the identification problem.

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 (2013), and Strzalecki (2014). The type structures in Kets (2010) and Heifetz and Kets (2013) capture finite-order beliefs about a primitive set of uncertainty, where the players may face uncertainty about the length of other players' 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 those papers, we do not include types with finite-order beliefs or beliefs about finite levels. This stems from the differences in the questions addressed by the papers. We are interested in understanding the extent to which behavior is consistent with high levels of reasoning. As a consequence, the ability 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 9E.)

2 Heuristic Treatment

Consider the game in Figure 2.1, where Player 1 is denoted by P1 (she) and Player 2 is denoted by P2 (he). We begin by applying the standard level-k solution concept to the game.

P1 \ P2	a_2	b_2	c_2	d_2
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 2.1: 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 2.2 describes the level-k behavior in four cases. In each case, 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 s_i is classified as level m for Pi.[2]

In each of these cases, there is no m so that d_i is classified as level m for Pi. If Pi has an anchor that assigns $0.5 : 0.5$ 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] As standard, we refer to the solution concept as "level-k." We use the index m to refer to a particular realization of k.

	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 2.2: Level-k

2. More generally:

Claim 2.1. *Suppose P1's and P2's anchors are given by (μ_1, μ_2). If d_i is level m for Pi, 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.[3] 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, there is no anchor μ_j under which c_j and d_j are both a best response.[4] 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 observes only 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 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 a "strategic" player.

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 understand better what goes into the statement that d_i is consistent with rationality and common belief of rationality. To do so, we

[3] When Pj is restricted to playing a strategy in $\{a_j, b_j, d_j\}$ (resp. $\{a_j, b_j, c_j\}$), d_i is dominated by a $0.25 : 0.75$ mixture on $a_i : b_i$ (resp. by a_i).

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

revisit a standard epistemic model, as applied to Figure 2.1. 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 is rational, we must describe what beliefs P1 holds about P2's play. After all, whether a strategy is a best response for P1 depends on her first-order belief. By similar logic, to specify whether P1 does or does not believe P2 is rational, we must describe P1's joint belief about P2's strategy and first-order belief, i.e., about his strategy and belief about P1'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, the set of types is:

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

Second, for each Pi, there is a belief map β_i, taking each type of Pi to a belief about the strategy-type pairs of Pj. In our example, the belief map is:

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

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), t_i assigns probability 1 to "Pj plays c_j and believes I play a_i." And so on. See Section 3.2 for the general case.

Now turn to rationality, belief in rationality, etc. Rationality is a property of a strategy-type pair: Whether a strategy is rational or irrational depends on the belief that a player holds, where the belief is specified by a type. For instance, the pair (a_i, t_i) is rational, since 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:

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

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$^{\text{st}}$-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 is 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 produced 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 1 to the other player Pj choosing c_j. Type t_1 has the first-order belief associated with P1's anchor μ_1 and type v_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 this anchor. Similarly, if P1's and P2's anchors (μ_1, μ_2) both assign probability 0.5 : 0.5 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 5.1 for a more complete argument.)

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 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 1 to Pj having a 1-type (i.e., their marginal beliefs on T_j assign probability 1 to T_j^1). And so on. Importantly, a level-k type structure is defined relative to a particular anchor and only induces hierarchies of beliefs consistent with that anchor. (See Proposition 5.1.) Thus, the type structure cannot induce all hierarchies of beliefs. (See Proposition 5.2.)

In principle, a level-k type structure can impose substantive assumptions about beliefs that go above and beyond the assumption that players' hierarchies are generated by the anchor. To rule out such substantive assumptions, we focus on, 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 induces that belief. Proposition 5.3 shows that there exists a complete level-k type structure that induces a rich set of beliefs consistent with the anchor. (See, also, Section 9A.)

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 6.1).** In any 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 to Figure 2.1. If we observe P1 play d_1, we cannot conclude that there is a bound m so that the behavior must reflect RmBR—importantly, we cannot draw this conclusion 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 7 discusses such additional auxiliary assumptions and the difficulty of verifying those assumptions in the data.

3 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.

Let I be a finite index set and $(\Omega_i : i \in I)$ be a collection of metrizable sets. 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., if $\omega_{-i} = (\omega_j : j \neq i)$, then $f_{-i}(\omega_{-i}) = (f_j(\omega_j) : j \neq i)$. 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 and, if f is continuous, $\underline{f}$ is continuous. (For measurability, see Friedenberg and Keisler, 2021, Lemma A.1; for continuity, see Aliprantis and Border, 2007, Theorem 14.14.)

Given a Borel set $\Phi \subseteq \Omega_1 \times \Omega_2$, write $\text{proj}_{\Omega_1} : \Phi \to \Omega_1$ for the projection mapping, i.e., the mapping with $\text{proj}_{\Omega_1}(\omega_1, \omega_2) = \omega_1$. If $\mu \in \Delta(\Phi)$, write $\text{marg}_{\Omega_1}\mu$ for the measure $\nu \in \Delta(\Omega_1)$ with $\nu(E_1) = \mu((\text{proj}_{\Omega_1})^{-1}(E_1))$ for each Borel $E_1 \subseteq \Omega_1$.

3.1 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 (that is, $|S_i| \geq 2$). 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 as implicit descriptions of the hierarchies of beliefs.

Definition 3.1. 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 t_i of player 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**.

3.2 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 first two ingredients are captured by G and the latter ingredient is captured by a type structure. The example in Section 2 indicates how types induce hierarchies of beliefs. In particular, each type $\tilde{t}_i$ induces a belief about the strategies about other players, given by $\text{marg}_{S_{-i}}\beta_i(\tilde{t}_i)$. For instance, type w_i's first-order belief assigns $0.5 : 0.5$ to $c_j : d_j$. Moreover, because each type has a joint belief about the strategy and type of the other player, each type has a joint belief about the strategy and first-order belief of the other player. For instance, type w_i assigns 0.5 to "Pj will play c_j and believes that I will play a_i" and 0.5 to "Pj will play d_j and assigns $0.5 : 0.5$ to me playing $c_i : d_i$." These joint beliefs constitute the type's second-order beliefs. And so on.

Begin by inductively describing the set of m^{th}-order beliefs of player i. Set $X_i^1 = S_{-i}$ and $H_i^1 = \Delta(X_i^1)$ and note both are compact metric spaces. Assume the sets X_i^m and H_i^m have been defined and are compact metric spaces. Set:

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

and $H_i^{m+1} = \Delta(X_i^{m+1})$. These, too, are compact metric spaces. (See Friedenberg, 2010, Lemma A1 and Remark A1.) The set X_i^m is player i's ***m*^*th*^-order space of uncertainty**. The set H_i^m is player i's set of **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 $\beta_i(t_i)$ on 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, viz., $\delta_i^2(t_i) = h_i^2$, is a joint belief about strategies and first-order beliefs: The probability that h_i^2 assigns to an event in $S_{-i} \times H_{-i}^1$ is 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 A.1 formally describes the maps $\delta_i^m : T_i \to H_i^m$. Given these maps, the map $\delta_i : T_i \to H_i^\infty$ 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 induces all hierarchies of beliefs.

Definition 3.2. Call the type structure $\mathcal{T}^*$ **hierarchy-complete** if, for each S-based type structure $\mathcal{T}$, $\mathcal{T}^*$ induces the hierarchies of beliefs induced by $\mathcal{T}$.[5]

To better understand the concept, consider a hierarchy-complete type structure $\mathcal{T}^* = (S_{-i}, T_i^*, \beta_i^* : i \in I)$, associated with hierarchy maps $\delta_i^* : T_i^* \to H_i^\infty$. If $\mathcal{T} = (S_{-i}, T_i, \beta_i : i \in I)$ is associated with hierarchy maps $\delta_i : T_i \to H_i^\infty$, then $\delta_i(T_i) \subseteq \delta_i^*(T_i^*)$. Thus, a hierarchy-complete type structure induces any hierarchy of beliefs that can be induced by any type structure. 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 hierarchy-complete.

[5]Some papers refer to this property as "universal." However, since the phrase universal is also used differently in the literature, we refer to this property as "hierarchy-complete."

4 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 ***i*'s anchor**.

Remark 4.1. 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 exceptions.) 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.

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. This section describes how the anchor restricts the partial hierarchies and, in turn, restricts the hierarchies of beliefs.

4.1 Hierarchies of Partial Beliefs

Much as types induce hierarchies of beliefs, they induce hierarchies of partial beliefs. To understand the difference between hierarchies of beliefs and hierarchies of partial beliefs, return to the example in Section 2. Each type $\tilde{t}_i$ induces a belief about the strategies about other players, given by $\text{marg}_{S_{-i}} \beta_i(\tilde{t}_i)$. This is both the type's first-order belief and first-order partial belief. That is, there is no distinction that arises at the first-order. While a type's second-order belief is a joint belief about strategies and first-order beliefs, the type's second-order partial belief is a belief only about first-order beliefs. So, type w_i's second-order partial belief assigns 0.5 to "Pj believes that I will play a_i" and 0.5 to "Pj assigns $0.5 : 0.5$ to me playing $c_i : d_i$." It does not include the statement that w_i assigns probability 0.5 to "Pj will play c_j and believes I will play a_i." As a consequence, it does not include information that, if w_i believes "Pj believes that I will play a_i," then w_i believes "Pj will play c_j." This is the sense in which this hierarchy is partial.

We now specify the hierarchies of partial beliefs. Set $P_i^1 = \Delta(S_{-i})$ and note that it is a compact metric space. Assume sets P_i^m have been defined and these are each compact metric spaces. Set $P_i^{m+1} = \Delta(P_{-i}^m)$ and note that it too is a compact metric space. Notice $P_i^2 = \Delta(\prod_{j \in I \setminus \{i\}} \Delta(S_{-j}))$ is the set of beliefs about first-order (partial) beliefs, while $H_i^2 = \Delta(\prod_{I \setminus \{i\}} (S_{-j} \times \Delta(S_{-j})))$ is the set of joint beliefs about strategies and first-order beliefs. More generally, the set P_i^m is player i's set of $\boldsymbol{m^{th}}$**-order partial 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 $\boldsymbol{\mu}_{-i} := (\mu_j : j \in I \setminus \{i\})$. 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 $p_{i,\boldsymbol{\mu}}^{m+1}(\{p_{-i,\boldsymbol{\mu}}^m\}) = 1$.

4.2 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, $p_i^2 = \text{marg}_{\prod_{j \neq i} \Delta(S_{-j})} h_i^2$, since a second-order partial belief simply provides information about beliefs over first-order partial beliefs and first-order partial beliefs are the 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 4.1. 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 $n \neq m$.)

5 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.

5.1 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 countable 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 5.1. 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. More generally, the $(m+1)$-types assign probability 1 to the m-types.

Example 5.1. To understand better what goes into a level-k type structure, return to the example of Section 2. That type structure is not a level-k type structure for any anchor $\boldsymbol{\mu} = (\mu_1, \mu_2)$. Indeed, 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_i so that u_i is an m_i-type. Observe that, if u_i is an m_i-type then u_{-i} is an $(m_i - 1)$-type. This implies that there must be some player i for which $u_i \in T_i^1$ and, therefore, $\mu_i(b_j) = 1$. As a consequence, u_i is the unique 1-type for that i.

Without loss of generality, suppose $T_1^1 = \{u_1\}$. Since, $T_1^1 = \{u_1\}$ and $T_2^2 \neq \emptyset$, it follows that $T_2^2 = \{u_2\}$. Iterating this argument,

$$T_1^{2m-1} = \{u_1\} \text{ and } T_2^{2m} = \{u_2\}$$

for each $m \geq 1$.

Observe, since t_2, u_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 1 to T_2^1, the set 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, there is no $m \geq 1$ so that w_1 is 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 induces only 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 5.1. *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 A.2 proves Proposition 5.1. 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 $p_{i,\boldsymbol{\mu}}^m$. So 1-types have first-order (partial) beliefs that coincide with the anchor; 2-types believe other players' first-order (partial) beliefs coincide with the anchor; and so on. This provides an interpretation of the m-types.

Because there are (always) hierarchies of beliefs that are inconsistent with the anchor, no level-k type structure can induce all hierarchies of beliefs.

Proposition 5.2. *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 hierarchy-complete.*

To understand the result, fix an anchor $\boldsymbol{\mu} = (\mu_i : i \in I)$ and a profile of first-order beliefs $\boldsymbol{\nu} = (\nu_i : i \in I)$ where each $\mu_i \neq \nu_i$. There is a profile of hierarchies of beliefs $(h_i : i \in I)$ at which player i's first-order beliefs are ν_i and this fact is commonly believed. But, that hierarchy cannot be induced by any type in any level-k type structure for $\boldsymbol{\mu}$: In a level-k type structure, an m-type must induce m^{th}-order partial beliefs consistent with the anchor $\boldsymbol{\mu}$. But, for each m, h_i induces m^{th}-order partial beliefs that are inconsistent with the anchor.

5.2 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 5.2. Consider a two-player game where each $S_i = \{\square_i, \diamond_i\}$. Suppose the anchor $\boldsymbol{\mu} = (\mu_1, \mu_2)$ is such that, for each i, $\mu_i(\square_{-i}) = \frac{2}{3}$. Consider a type structure $\mathcal{T}$ with the following properties: Set $T_1 = T_2 = \mathbb{N}^+$. Take each $\beta_i(1)$ so that $\beta_i(1)(\square_{-i}, 2) = \frac{2}{3}$ and $\beta_i(1)(\diamond_{-i}, 3) = \frac{1}{3}$. For $m \geq 2$, take

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

and

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

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

Example 5.3. Consider a two-player game where each $S_i = \{\square_i, \diamond_i\}$. Suppose the anchor $\boldsymbol{\mu} = (\mu_1, \mu_2)$ is such that, for each i, $\mu_i(\square_{-i}) = \frac{2}{3}$. Consider a type structure $\mathcal{T}$ with the following properties: Set $T_1 = T_2 = \mathbb{N}^+$. Take each $\beta_i(1)$ so that $\beta_i(1)(\square_{-i}, 2) = \frac{2}{3}$ and $\beta_i(1)(\diamond_{-i}, 3) = \frac{1}{3}$. For $m \geq 2$, take

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

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

Examples 5.2-5.3 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 5.1.

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 5.2, each such m-type has a degenerate belief, assigning probability 1 to either of $\square_{-i}$ or $\diamond_{-i}$; in Example 5.3, each such m-type has a non-degenerate belief, assigning $\frac{1}{2} : \frac{1}{2}$ to $\square_{-i} : \diamond_{-i}$. Thus, for each type $m \geq 2$ in Example 5.2, there is no type $n \geq 1$ in Example 5.3 that induces the same first-order beliefs, a fortiori the same hierarchies of beliefs. And conversely, with Example 5.3 and Example 5.2 reversed. Moreover, the 1-types induce distinct second-order beliefs. In Example 5.2, 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 Example 5.3, type 1 assigns zero probability to that same event.

5.3 Complete Level-k Type Structures

Proposition 5.1 says that a level-k type structure imposes the substantive assumption that the hierarchies are induced by the anchor. But, Examples 5.2-5.3 illustrated that there may be multiple level-k type structures, associated with the same anchor, which induce different hierarchies of beliefs. To understand why this arises, note that, in Examples 5.2-5.3 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 beyond the substantive assumptions imposed by the anchor. We will be interested in type structures that do not impose these exogenous restrictions on beliefs (or, at least, minimize such exogenous restrictions).

Definition 5.2. 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.

Thus, $\mathcal{T}$ 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 5.3. *Fix some $\boldsymbol{\mu} = (\mu_i : i \in I)$. 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 5.3 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.[6] Thus, there are no restrictions on the beliefs of 1-types aside from the requirement that their first-order beliefs coincide with the anchor (and that they have higher-order beliefs consistent with the type structure). 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.[7] Section 9A provides an example and a broader discussion of missing hierarchies. Section 9B discusses why any missing hierarchies are immaterial from the perspective of the inference problem.

6 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.[8] Reasoning about rationality will be captured by the epistemic conditions of *rationality and m^{th}-order belief of rationality.*

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

[7] It is also the case that the structure constructed in Proposition 5.3 is not type-complete, in the sense that the belief maps are not onto.

[8] 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.

6.1 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.

Given some $\nu_i \in \Delta(S_{-i})$, write $\mathbb{BR}_i[\nu_i]$ for the set of strategies $s_i \in S_i$ with $\pi_i(s_i, \nu_i) \geq \pi_i(r_i, \nu_i)$ for all $r_i \in S_i$.

Definition 6.1. 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 6.2. 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 6.3. The set of states at which there is **rationality and m^{th}-order belief of rationality** (**RmBR**) 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 $R^\infty = \prod_{i \in I} R_i^\infty$.

6.2 The Unrestricted Inference Problem

The unrestricted inference problem is not our focus of 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. Less formally, 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 $s_i \in \mathbb{BR}_i[\nu_i]$ and $\nu_i(S_{-i}^m) = 1$. 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 6.1 (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 hierarchy-complete, for each $m \geq 1$, $\text{proj}_S R^m = S^m$.

(iii) If $\mathcal{T}$ is hierarchy-complete, $\text{proj}_S R^\infty = S^\infty$.

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

To understand how this 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 hierarchy-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, s_i is consistent with unbounded reasoning about rationality. In particular, in any hierarchy-complete structure, there is a type t_i so that (s_i, t_i) satisfies RCBR (part (iii)).

6.3 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 6.1. *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-rationalizable 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 in 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.

Example 6.1. Consider the game in Figure 6.1. Note, for each player Pi, b_i is a best response under any $\nu_i \in \Delta(\{a_{-i}, b_{-i}\})$, but a_i is a best response under $\nu_i \in \Delta(\{a_{-i}, b_{-i}\})$ if and only if $\nu_i(a_i) = 1$. The entire strategy set is rationalizable.

Consider now a level-k type structure for $\boldsymbol{\mu}$, where each $\mu_i(a_{-i}) < 1$, and $\mathcal{T} = (S_{-i}, T_i, \beta_i : i \in I)$. We will argue that there is no type $t_i \in T_i$ with $(a_i, t_i) \in R_i^\infty$.

		P2	
		a_2	b_2
P1	a_1	1, 1	0, 1
	b_1	1, 0	1, 1

Figure 6.1: RCBR in a Level-k Type Structure

Suppose, contra hypothesis, that there is an $(a_i, t_i) \in R_i^{\infty}$. Then, t_i's first-order belief $\delta_i^1(t_i)$ must assign probability 1 to a_{-i}. Write $h_{i,a}^1 = \delta_i^1(t_i)$. Type t_i's second-order belief $\delta_i^2(t_i)$ must assign probability one to $(a_{-i}, h_{-i,a}^1)$. Inductively, for each $m \geq 1$, type t_i's $(m+1)^{th}$-order belief $\delta_i^{m+1}(t_i)$ must assign probability one to $(a_{-i}, h_{-i,a}^1, \ldots, h_{-i,a}^m)$.

Now, observe that, since $\mathcal{T}$ is a level-k type structure for $\boldsymbol{\mu}$, there is some $\ell \geq 1$ so that $t_i \in T_i^{\ell}$. Thus, type t_i has an ℓ^{th}-order partial belief that corresponds to the anchor. But, this contradicts the fact that $\delta_i^{\ell}(t_i) = h_{i,a}^{\ell}$.[9]

6.4 Proof of Theorem 6.1

We now turn to prove Theorem 6.1. Part (i) is an implication of Proposition 6.1'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 $\operatorname{marg}_{S_{-i}} \beta_i(t_i^2) = \nu_i$. 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 $\operatorname{marg}_{S_{-i}} \beta_i(t_i^{m+2}) = \nu_i$, $(s_i, t_i^{m+2}) \in R_i^1$. Moreover, for each $n \leq m$, R_{-i}^n is Borel (Lemma B.3) and $\operatorname{Supp} \beta_i(t_i^{m+2}) \subseteq R_{-i}^m \subseteq R_{-i}^n$. So, t_i^{m+2} believes R_{-i}^n for each $n \leq m$. As such, $(s_i, t_i^{m+2}) \in R_i^{m+1}$.

7 The Level-k Inference Problem

Theorem 6.1 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.

7.1 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.

[9]This paragraph is formalized in the proof of Proposition 5.2.

Definition 7.1. 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) $\nu_i(L_{-i}^m) = 1$.

Say a strategy is **level** m (**for** $\boldsymbol{\mu}$) if $s_i \in L_i^m$. Call the set L_i^m as ***i*'s level** m **behavior** (**for** $\boldsymbol{\mu}$) and call the set $L^m = \prod_{i \in I} L_i^m$ 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 1 to the level-1 behavior of other players. And so on.

Remark 7.1. Our definition allows for the fact that the sets L_i^m 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 (Ledoux, 1981; 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 point further in Section 9D.

7.2 Epistemic Foundations for Level-k

Theorem 7.1. *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* $\mathcal{C}_i = \{T_i^m : m = 1, 2, \ldots\}$ *satisfying conditions (i)-(ii) of Definition 5.1 (resp. (i)-(ii)-(iii) of Definition 5.2, if* $\mathcal{T}$ *is a complete level-*k *type structure).*

(i) For each $m \geq 1$*,* $\operatorname{proj}_{S_i}(R_i^m \cap (S_i \times T_i^m)) \subseteq L_i^m$*.*

(ii) If $\mathcal{T}$ *is a complete level-*k *type structure for* $\boldsymbol{\mu}$*, for each* $m \geq 1$*,* $\operatorname{proj}_{S_i}(R_i^m \cap (S_i \times T_i^m)) = L_i^m$*.*

Much like Theorem 6.1, Theorem 7.1 fixes a level-k type structure for $\boldsymbol{\mu}$. Refer to Figure 7.1. Whereas Theorem 6.1 focused on the behavioral implications of R$(m-1)$BR (Figure 7.1(a)), Theorem 7.1 focuses on the behavioral implications of R$(m-1)$BR for *only the* m*-types* (Figure 7.1(b)). 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 for $\boldsymbol{\mu}$ if and only if there is a 1-type t_i so that (s_i, t_i) is rational. (See Proposition B.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.

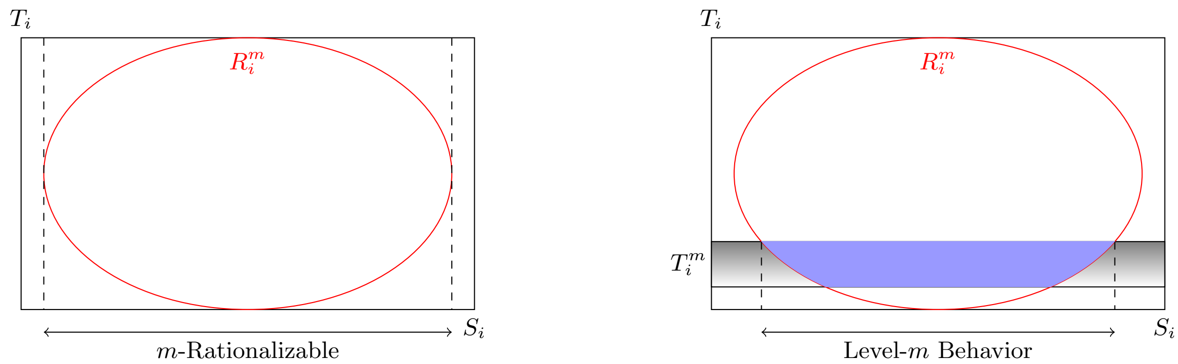

(a) Projection of R$(m-1)$BR

(b) Projection of R$(m-1)$BR and m-Types

Figure 7.1: Theorem 6.1 vs Theorem 7.1

Example 7.1. Refer to the game in Figure 7.2. Consider an anchor $\boldsymbol{\mu} = (\mu_1, \mu_2)$ with each $\mu_i(c_{-i}) = 1$. Observe that

$$L_i^m = \{a_i, b_i\} = S_i^m$$

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 R$(m-1)$BR, 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.

		P2 a_2	b_2	c_2
	a_1	1, 1	0, 0	1, -1
P1	b_1	0, 0	1, 1	1, -1
	c_1	-1, 1	-1, 1	-1, -1

Figure 7.2

The type structure has type sets $T_i = \{t_i, v_i\} \times \mathbb{N}$ and belief maps that satisfy the following: First, $\beta_i(t_i, 1)(c_{-i}, (t_{-i}, 2)) = \beta_i(v_i, 1)(c_{-i}, (v_{-i}, 2)) = 1$ and $\beta_i(t_i, 2)(a_{-i}, (t_{-i}, 1)) = \beta_i(v_i, 2)(a_{-i}, (v_{-i}, 1)) = 1$. Second, $\beta_i(t_i, 3)(a_{-i}, (t_{-i}, 2)) = \beta_i(v_i, 3)(c_{-i}, (v_{-i}, 2)) = 1$. Third, for each $m \geq 4$, $\beta_i(t_i, m)(a_{-i}, (t_{-i}, m-1)) = \beta_i(v_i, m)(b_{-i}, (v_{-i}, m-1)) = 1$.

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, $\text{proj}_{S_i} R_i^m = \{a_i, b_i\}$. However, $\text{proj}_{S_i}(R_i^2 \cap T_i^2) = \{a_i\} \subsetneq L_i^m$.

Example 7.1 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 6.1 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 7.1 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 strategy 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.

7.3 Identifying Levels of Reasoning about Rationality

Suppose the analyst observes a player choose some strategy s_i^* so that (i) s_i^* is level $m \geq 1$ (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 7.1. To do so, we use the following (well-known) fact: If s_i^* is level m then s_i^* is m-rationalizable. (See Lemma B.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 6.1.

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 6.1 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 7.1 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 7.1.) 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 partial n^{th}-order beliefs consistent with the anchor, then they reason according to*

$R(n-1)BR$.[10] This is an assumption that a player's n^{th}-order beliefs determine the extent to which the player reasons about rationality—an assumption that would be hard to falsify or verify in practice.

Remark 7.2. This section asked: What can the observer infer about how a player reasons *about rationality*, if they observe a strategy that is level m but not level n for any $n > m$. It is worth emphasizing that if behavior is consistent with reasoning about rationality beyond level m, then it is also consistent with other forms of reasoning beyond level m. For instance, it would be consistent with unbounded interactive reasoning—i.e., statements of the form "I think, you think, etc…." In some games (including games of substantive interest), it is also consistent with reasoning about irrationality beyond level m.

8 Applications

The level-k model has served as an important lens through which experimentalists have evaluated bounded reasoning. This section revisits three prominent games—each studied broadly in the experimental literature—through the lens of Theorem 6.1. It argues that, even when the hierarchies are generated by the anchors that the experimental literature has focused on, the observed behavior is often consistent with higher levels of reasoning about rationality.

8.1 Beauty Contest

The beauty contest game was initially studied by Ledoux (1981) and Nagel (1995). Each of $|I| \geq 3$ players compete for a prize, whose value is 1. They do so by simultaneously choosing a number $s_i \in \{\underline{x}, \underline{x}+\Delta, \ldots, \underline{x}+\kappa\Delta\}$, where $\underline{x} \geq 0$, $\Delta > 0$ and $\kappa \geq 1$ is an integer. Player i is a winner if her chosen number best matches a fraction $p \in (0,1)$ of the average strategy chosen. That is, given a strategy profile $s = (s_j : j \in I)$, the set of winners is

$$W(s_j : j \in I) = \{i \in I : |s_i - \frac{p}{|I|}\sum_{j \in I} s_j| \leq |s_\ell - \frac{p}{|I|}\sum_{j \in I} s_j| \text{ for each } \ell \in I\}.$$

Player i's payoff function is given by

$$\pi_i(s) = \begin{cases} \frac{1}{|W(s)|} & \text{if } i \in W(s) \\ 0 & \text{otherwise.} \end{cases}$$

As an example, suppose $|I| = 3$, the set of strategies is $\{0,1,2,3,4,5\}$, and $p = \frac{2}{3}$. Consider an anchor $\boldsymbol{\mu}$, where each μ_i is uniform. Given this anchor, player i's expected payoff is maximized by choosing 1.[11] Thus, $L_i^1 = \{1\}$ and, for each $m \geq 2$, $L_i^m = \{0\}$. Yet, even if players' hierarchies of beliefs are generated by the uniform anchor, observing $s_i \in \{1,2,3,4\}$ is consistent with a higher bound on reasoning about rationality. To see this, note that, for each $m \geq 1$,

$$S_i^m = \begin{cases} \{0,\ldots,5-m\} & \text{if } 5 \geq m \\ \{0\} \text{ if } m \geq 6. \end{cases}$$

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

[11] The expected payoff from 1 is approximately .65, which is higher than the expected payoffs of the other strategies: for 0 it is .31, for 2 it is .62, for 3 it is .31, for 4 it is .09, and for 5 it is .01.

So, even if players beliefs are generated by the anchor, an observation of $s_i = 2$ (resp. $s_i = 1$) is consistent with R1BR (resp. R2BR).

The particular parameterization of the game is important for understanding the extent to which behavior in the beauty contest can indicate bounded reasoning. (Again, even if the beliefs are generated by an anchor.) To see this, set

$$\hat{x}(I, \Delta, p) = \frac{\Delta(|I| - 2p)}{2|I|(1-p)}.$$

Seel and Tsakas (2017) show that, for each $m \geq 1$, the m-rationalizable strategies take the following form:

$$S_i^m = \begin{cases} \{\underline{x}, \ldots, \underline{x} + (\kappa - m)\Delta\} & \text{if } \kappa \geq m \text{ and } \underline{x} + (\kappa - m + 1)\Delta > \hat{x}(I, \Delta, p) \\ S_i^{m-1} & \text{otherwise.} \end{cases}$$

To understand the characterization more fully, focus on two extremes. First, suppose $\underline{x} + \kappa\Delta \leq \hat{x}(I, \Delta, p)$. (This will, for instance, be the case when the actions are $\{0, 1, 2, 3, 4, 5\}$, $p = .95$, and $|I| \geq 4$.) In this case, any strategy is consistent with R1BR, R2BR, R3BR, etc..., in a complete level-k type structure (for $\boldsymbol{\mu}$). Second, suppose that $\underline{x} + \Delta > \hat{x}(I, \Delta, p)$. In that case, player i always wins the full prize, if the other players choose the same action and player i undercuts the other players by Δ. So, all but the highest strategy is 1-rationalizable; all but the top two strategies are 2-rationalizable, etc. As such, in a complete level-k type structure (for $\boldsymbol{\mu}$), a strategy $\underline{x} + \lambda\Delta$ may be consistent with R$(m-1)$BR even though there is no $n \geq m$ so that $\underline{x} + \lambda\Delta$ is categorized as level n (for $\boldsymbol{\mu}$).

In experiments, the grid size is often small relative to p and $|I|$. More concretely, we can think of a player's strategy set as a finite subset of $[0, \infty)$, with a lowest strategy of $\underline{x}$ and a highest strategy of $\overline{x} > \underline{x}$. Then, $\kappa\Delta$ represents the bound on the available strategies, i.e., $\kappa\Delta = \overline{x} - \underline{x}$. Experiments typically set $\overline{x} - \underline{x}$ and implement a fine grid; this is captured by Δ small and κ high. (See the discussion in Section 9F for a formalization.) With a fine grid, rationalizability results in eliminating strategies "slowly." For example, take $|I| = 3$, $p = \frac{2}{3}$, and $[\underline{x}, \overline{x}] = [0, 1]$ with a grid given by $(\Delta_n, \kappa_n) = (\frac{1}{n}, n)$. When $n = 2$, $s_i = \frac{1}{2}$ is not 2-rationalizable. But, $s_i = \frac{1}{2}$ is 2-rationalizable when $n = 100$; in fact, it is 49-rationalizable. Contrast this with the level-k classification, where the anchor is taken either to be uniform or to assign probability 1 to the arithmetic mean.[12] When the anchor assigns probability 1 to the arithmetic mean y, behavior above (the grid member closest to) $p \times y$ is categorized as "irrational" or "unsophisticated" and behavior just below is categorized as reflecting exactly 1 level of reasoning. In the example, behavior above $\frac{1}{2}$ is viewed as irrational (irrespective of the grid); in fact, when $n = 100$, behavior above .35 is viewed as irrational. Notice, when the grid is fine, behavior viewed as irrational or reflecting exactly 1 level of reasoning is, in fact, consistent with high levels of reasoning about rationality—even when the hierarchies are generated by the same anchor. Section 9F discusses an additional reason that level-k classification may underestimate the extent to which there is bounded reasoning, in the fine-grid beauty contest.

[12] These are different in practice. See Breitmoser (2012).

8.2 Guessing Games

An important variant of the beauty contest is a guessing game, first studied in Costa-Gomes and Crawford (2006). These games differs from the beauty contest in that their payoff functions are now

$$\pi_i(s_j : j \in I) = B - \left(s_i - \frac{p}{|I|-1} \sum_{j \in I \setminus \{i\}} s_j \right)^2,$$

where $B \geq 0$ and $p \in (0,1)$.[13] So, each player i seeks to minimize the distance to their expectation of the target.

There are two strategic differences from the beauty contest. First, the target only depends on the average behavior of the other players and is not impacted by a player's own choice. Second, now, player i has a *direct* incentive to match the target; because the target depends on p, player i may have an incentive to choose a strategy that is significantly lower than their expectation of the average strategy. (See Lemma C.4.) By contrast, in the beauty contest, player i only has an incentive to be the player closest to the target; this is an *indirect* incentive to match the target, which can, in principle, be met with higher strategies.

To see the implication of these new incentives, suppose $|I| = 3$, the set of strategies is $\{0,1,2,3,4,5\}$, and $p = \frac{2}{3}$. Consider an anchor $\boldsymbol{\mu}$, where each μ_i is uniform. Now, $L_i^1 = \{2\}$ and, for each $m \geq 2$, $L_i^m = \{1\}$. (Notice, $L_i^3 = \{1\}$ precisely because a player's own action doesn't impact the target.) Yet, even if players' hierarchies of beliefs are generated by the uniform anchor, observing $s_i \in \{0,2,3\}$ is consistent with a higher bound on reasoning about rationality. To see this, note that, for each $m \geq 1$,

$$S_i^m = \begin{cases} \{0,1,2,3\} & \text{if } m = 1 \\ \{0,1,2\} & \text{if } m = 2 \\ \{0,1\} & \text{if } m \geq 3. \end{cases}$$

So, even if players beliefs are generated by the anchor, an observation of $s_i = 2$ is consistent with R1BR and an observation of $s_i = 0$ is consistent with RmBR for all $m \geq 1$.

The m-rationalizable strategies are given by

$$S_i^m = \{\underline{x}, \ldots, \underline{x} + J(m)\Delta\},$$

where $J : 0,1,2,\ldots \to \{0,\ldots,\kappa\}$ is a weakly decreasing function satisfying the following criteria:

(i) $J(0) = \kappa$;

(ii) $J(m+1) < J(m)$ if $J(m) \geq 1$ and $\underline{x} + J(m)\Delta > \tilde{x}(\Delta,p) := \frac{\Delta}{2(1-p)}$; and

(iii) $J(m+1) = J(m)$ otherwise.

As the example indicated, $J(m+1)$ can be strictly less than $J(m) - 1$, i.e., on any given round more than one strategy can be eliminated. (In the specific case where $\underline{x} + \kappa\Delta \in (\tilde{x}(\Delta,p), 3\tilde{x}(\Delta,p)]$, rationalizability begins by eliminating the highest strategy on each round, until some round m where the m-rationalizable strategies stop shrinking. See Appendix C.)

[13] There are variants based on the distance between s_i and $p/(|I|-1) \sum_{j \in I \setminus \{i\}} s_j$, as measured by the absolute value. A similar discussion applies to those variants.

In experiments, the grid size is small relative to p. In that case, rationalizability involves eliminating more than simply the highest strategy. But, the anchor is typically taken to be uniform or to assign probability 1 to the arithmetic mean. For these anchors, the highest 1-rationalizable strategy is typically higher than the highest level 1 strategy, despite the fine grid. Lemma C.4 points to why. For instance, take $[\underline{x}, \overline{x}] = [0, 1]$ with a grid given by $(\Delta_n, \kappa_n) = (\frac{1}{n}, n)$; again, a fine grid corresponds to $n \geq 1$ large. When the anchor is uniform, a player's best response is to choose the strategy (in the grid) closest to $\frac{p}{2}$. But, when the grid is fine, there are 1-rationalizable strategies in $(\frac{p}{2}, p)$; in fact, when $p > \frac{1}{2}$ and the grid is fine, there are 2-rationalizable strategies in $(\frac{p}{2}, p^2)$, etc. So, even when hierarchies are generated by the uniform anchor, behavior classified as level 0, level 1, etc., may be consistent with higher rounds of reasoning about rationality.

8.3 11-20 Game

This game was initially studied in Arad and Rubinstein (2012). Two players simultaneously choose a number in $\{11, 12, \ldots, 20\}$. In the original version of the game (Arad and Rubinstein, 2012), payoff functions are given by

$$\pi_i(s_i, s_{-i}) = \begin{cases} s_i + 20 & \text{if } s_{-i} = s_i + 1 \\ s_i & \text{otherwise.} \end{cases}$$

Alaoui and Penta (2016) propose a version where the payoff functions are given by

$$\pi_i(s_i, s_{-i}) = \begin{cases} s_i + 20 & \text{if } s_{-i} = s_i + 1 \\ s_i + 10 & \text{if } s_{-i} = s_i \\ s_i & \text{otherwise.} \end{cases}$$

Notice, in both versions, each strategy $s_i \in \{11, 12, \ldots, 19\}$ is a best response to $s_{-i} = s_i + 1$. In the original version (Arad and Rubinstein, 2012), $s_i = 20$ is a best response to $s_{-i} = 11$; in the revised version (Alaoui and Penta, 2016), $s_i = 11$ is a best response to $s_{-i} = 11$. Note that these differences are a deliberate feature of the design, both in Arad and Rubinstein and in Alaoui and Penta.

The two games are different from the perspective of the level-k concept. To see this, consider an anchor with each $\mu_i(20) = 1$. (See point (i) on page 3563 of Arad and Rubinstein on why this is a natural anchor.) In both games, for each $m = 1, \ldots, 9$, $L_i^m = \{20 - m\}$. In the original version of the game $L_i^{10} = \{20\}$ and the levels cycle. As a consequence, the initial version of the game, a strategy classified as level $m \geq 1$ if and only if it is classified as level $\ell + m$ for $\ell \in \{10, 20, 30, \ldots\}$. But, in the revised version of the game, $L_i^m = \{11\}$ for each $m \geq 10$.

The two games do not differ from the perspective of rationalizability. In particular, an important and deliberate feature of the design is that the entire strategy set is rationalizable.[14] (In Alaoui and Penta's version, 20 is a best response under a belief that assigns $\frac{19}{20} : \frac{1}{20}$ to $11 : 20$.) As a consequence, even in players' hierarchies are generated by an anchor that assigns probability 1 to 20, every strategy is consistent with RmBR for all $m \geq 1$.

For the original version of the game, the classification from the level-k model cannot overestimate a bound on reasoning: Any strategy is consistent with unbounded levels of reasoning according to the level k

[14]See point (vi) on page 3563 in Arad and Rubinstein on why this is desirable.

model. (For this reason, Arad and Rubinstein are careful to only draw conclusions about a bound when the subjects' own explanations allowed them to do so.) The conclusion is different for the revised version of the game. When the anchor assigns probability 1 to 20, the level-k model suggests the strategy $20 - m \geq 12$ has a bound of at most m. Nevertheless, all strategies are consistent with unbounded reasoning about rationality, even when the hierarchies are generated by the anchor.

9 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 9.1. Consider a two-player game where each $S_i = \{\square_i, \diamond_i\}$. For each player i, there is a hierarchy of beliefs $h_{i,\square} = (h^1_{i,\square}, h^2_{i,\square}, \ldots)$ where it is commonly believed that the other player chooses $\square_{-i}$: So, $h^1_{i,\square}(\square_{-i}) = 1$ and $h^{m+1}_{i,\square}(\square_{-i}, \ldots, h^m_{-i,\square}) = 1$. Also, for each player i, there is a hierarchy of beliefs $h_i = (h^1_i, h^2_i, \ldots)$ with $h^1_i(\square_{-i}) = \frac{2}{3}$, $h^{m+1}_i(\square_{-i}, \ldots, h^m_{-i,\square}) = \frac{2}{3}$, and $h^{m+1}_i(\diamond_{-i}, \ldots, h^m_{-i,\square}) = \frac{1}{3}$. (So, h^2_i assigns probability $\frac{2}{3}$ to "the other player plays $\square_{-i}$ and believes I play $\square_i$" and probability $\frac{1}{3}$ 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 $h^1_i = p^1_{i,\boldsymbol{\mu}}$. However, there is no level-k type structure (including a complete level-k type structure) that induces the hierarchy $h_i = (h^1_i, h^2_i, \ldots)$. 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 $\mathcal{C}_i = \{T^m_i : m = 1, 2, \ldots\}$ be a Borel cover so that $(\mathcal{C}_i : i \in I)$ jointly satisfy conditions (i)-(ii) of Definition 5.1. 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 $h^1_{i,\square}$. Since the 2-types must assign probability 1 to 1-types, this implies that the 2-types have second-order beliefs distinct from $h^1_{i,\square}$. 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.[15] As a consequence:

Proposition 9.1. *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

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

of Definition 5.2 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 5.3 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 A.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 6.1(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 with R$(m-1)$BR strategy in a complete level-k type structure (Theorem 6.1(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 5.1) 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 5.2) 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 5.3 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 7.1 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 7.1, some papers instead assume that players have a uniform belief about best responses. This imposes a secondary exogenous restriction on beliefs—one that depends on iterative best responses. This additional restriction only serves to reinforce the message of the paper: It gives a level-k bound that is lower than that suggested by Definition 7.1. As such, this choice may suggest lower levels of reasoning about rationality than is consistent with the data.

Theorem 7.1 can be seen as providing foundations for this level-k solution concept, as specified by Definition 7.1. 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 the opposite direction: It derives first-order beliefs based on best responses (to other beliefs).

E. Foundations for Level-k Theorem 7.1 provides epistemic foundations for the level-k solution concept. These foundations are quite different from foundations for other solution concepts: 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. By contrast, the foundations here 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. It is this property that leads leads to the difficulty with identification discussed in Section 7.3.

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 ...." Often, the level-k solution concept is motivated by a stipulation that players have a limited ability to engage in such sentences. Theorem 7.1 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.

Our framework is expressive beyond what might appear to be needed for certain results. But there is no sense that the additional expressiveness interferes with the conclusion of the results. The key is that the epistemic conditions of RmBR depend only on the $(m+1)^{\text{th}}$-order beliefs.[16] Importantly, this conclusion remains true even if the epistemic model contains types that consider the possibility that other players "do not reason" (as in Heifetz and Kets, 2013 or Kets, 2010). See Appendix D for a formal statement.

F. Fine Grid: Beauty Contest Return to the beauty contest. One might have thought that a fine grid uses a special structure that brings the sets of level-k and rationalizable strategies into close (or even exact) agreement. If this were to happen, it would weaken the import of our results. But this is not the case. Section 8.1 already pointed to the fact that a fine grid can exacerbate the extent to which the level-k classification underestimates reasoning. Here we see a second issue: In any fine-grid beauty contest game, there may multiple rationalizable strategies. So, even when hierarchies are generated by a given anchor, there may be multiple strategies that are consistent with unbounded reasoning about rationality.

Example 9.2. Suppose $|I| = 10$ and $p = .9$. Choose the strategy set so that $[\underline{x}, \overline{x}] = [0, 1]$ with a grid given by $(\Delta_n, \kappa_n) = (\frac{1}{n}, n)$. Then, the m-rationalizable strategies will be a strict subset of the $(m-1)$-rationalizable strategies if and only if $\kappa_n - m \geq 4$. When Δ_n is small, κ_n is large. Thus, we can find an m so that $4 > \kappa_n - m \geq 1$. In that case, the set of rationalizable strategy profiles is not a singleton.

In the limiting case, where each $S_i = [\underline{x}, \overline{x}]$, the set $[\underline{x}, \overline{x})$ is contained in the rationalizable strategy set for i. (This is irrespective of I and p.) See Appendix D.4.

To understand why this occurs, consider a sequence $((\kappa_n, \Delta_n) : n \geq 1)$ so that, for each $n \geq 1$, $\underline{x} + \kappa_n \Delta_n = \overline{x}$ and $\lim_{n\to\infty} \Delta_n = 0$. Then, for each $\varepsilon > 0$, there exists $N(\varepsilon) \geq 1$ so that $\Delta_n < \varepsilon$ for all $n \geq N(\varepsilon)$. A **ε-fine grid** represents a grid with $(\kappa, \Delta) = (\kappa_n, \Delta_n)$ for some $n \geq N(\varepsilon)$. In the beauty contest, for any $m \geq 1$, there is a ε-fine grid so that the set of m-rationalizable strategies is strictly contained in the set of $(m-1)$-rationalizable strategies. (See Appendix D.4.) But this stops short of saying that there is a $\hat{\varepsilon} > 0$ so that, for any $\varepsilon > \hat{\varepsilon}$, there is a unique rationalizable strategy for a game with a ε-fine grid. In fact, that may not be the case, as indicated by the example.

[16] This can be seen by recasting standard results in hierarchy spaces.

10. Conclusion

This paper takes seriously the idea that players' hierarchies of beliefs are shaped by an anchor, the key assumption associated with level-k models. Toward that end, it focuses on type structures that capture the substantive assumptions that hierarchies are induced by an anchor. In a sense, complete level-k type structures don't impose (needless) auxiliary assumptions on beliefs, above and beyond the requirement that hierarchies are induced by an anchor. So, by analyzing RmBR in a complete level-k type structure, we capture the reasoning (or "cognitive") process typically associated with the the level-k model. Nonetheless, Theorem 6.1 shows that this reasoning process has less predictive power than the level-k solution concept; it has the same predictive power as rationalizability.

Theorem 7.1 points to a reasoning process that has the predictive power of the level-k solution concept. But, the same result points to a new identification assumption—one that challenges the ability to infer a "level of reasoning" from the fact that behavior is classified as some level m. Do there exist alternate assumptions about beliefs—assumptions that are testable—which would allow the researcher to infer a subject's "level of reasoning" from the fact that behavior is classified as level m? Note, such assumptions would go above and beyond that discussed in the level-k literature. As such, this is a question for future research.

Appendix A Proofs for Sections 4-5

A.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 $\rho_i^m : S_{-i} \times T_{-i} \to X_i^m$ and $\delta_i^m : T_i \to H_i^m$. First, set $\rho_i^1 = \text{proj}_{S_{-i}}$ and $\delta_i^1 = \underline{\rho}_i^1 \circ \beta_i$. Note, ρ_i^1 is measurable and so $\underline{\rho}_i^1$ is measurable. From this and the fact that β_i is measurable, δ_i^1 is measurable.

Now, assume the measurable maps $\rho_i^m : S_{-i} \times T_{-i} \to X_i^m$ and $\delta_i^m : T_i \to H_i^m$ have been defined. Set

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

Note, since ρ_i^m and δ_{-i}^m are measurable, so is ρ_i^{m+1}. Then set $\delta_i^{m+1} = \underline{\rho}_i^{m+1} \circ \beta_i$. Since ρ_i^{m+1} is measurable and so $\underline{\rho}_i^{m+1}$ is measurable. From this and the fact that β_i is measurable, δ_i^{m+1} is measurable.

The following standard lemmata will be of use.

Lemma A.1. *For each $t_i \in T_i$, $\delta_i^1(t_i) = \text{marg}_{S_{-i}} \beta_i(t_i)$.*

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 A.2. *For each $m \geq 1$, $\delta_i^m(t_i) = \text{marg}_{X_i^m} \delta_i^{m+1}(t_i)$.*

Proof. Fix some Borel $E_i^m \subseteq X_i^m$. Note that

$$\delta_i^m(t_i)(E_i^m) = \beta_i(t_i)((\rho_i^m)^{-1}(E_i^m)) = \beta_i(t_i)((\rho_i^{m+1})^{-1}(E_i^m \times H_{-i}^m)) = \delta_i^{m+1}(t_i)(E_i^m \times H_{-i}^m),$$

as required. ■

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

Proof. For $m = 1$, this is immediate. Assume the statement is true for $m \geq 2$, 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. ■

A.2 Proof of Proposition 5.1

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 5.1. The following Lemma will establish Proposition 5.1.

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

Proof. The case of $m = 1$ is immediate. Assume the claim holds for $m \geq 2$. 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 $\eta_i^{m+1}(h_i^{m+1}) = p_{i,\boldsymbol{\mu}}^{m+1}$.

To see this, fix

$$F_{-i}^m = \{p_{-i,\boldsymbol{\mu}}^m\} \subseteq P_{-i}^m = Y_i^{m+1}$$

and observe that F_{-i}^m is measurable. Note, $\eta_i^{m+1}(h_i^{m+1}) = p_{i,\boldsymbol{\mu}}^{m+1}$ if and only if $\eta_i^{m+1}(h_i^{m+1})(F_{-i}^m) = 1$. Observe

$$\begin{aligned}\eta_i^{m+1}(h_i^{m+1})(F_{-i}^m) &= h_i^{m+1}((\hat{\eta}_i^{m+1})^{-1}(F_{-i}^m)) \\ &= h_i^{m+1}\left(\{(x_{-i}^m, h_{-i}^m) \in X_i^{m+1} : \eta_i^m(h_{-i}^m) = p_{-i,\boldsymbol{\mu}}^m\}\right).\end{aligned}$$

Thus,

$$\eta_i^{m+1}(h_i^{m+1})(F_{-i}^m) = h_i^{m+1}(E_i^{m+1}),$$

where

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

As such, it suffices to show that $h_i^{m+1}(E_i^{m+1}) = 1$.

To show this, first observe that

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

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

Now note that

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

where the inequality follows from the fact that $S_{-i} \times T_{-i}^m \subseteq (\rho_i^{m+1})^{-1}(E_i^{m+1})$ and the last equality follows from the fact that $t_i \in T_i^{m+1}$. From this, $h_i^{m+1}(E_i^{m+1}) = 1$ as desired. ■

A.3 Proof of Proposition 5.2

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 5.1. We will show that $\mathcal{T}$ is not hierarchy-complete. To do so, we construct an alternate type structure and show that it induces hierarchies of beliefs that are not induced by $\mathcal{T}$.

Consider an alternate type structure $\mathcal{T}^* = (S_{-i}, T_i^*, \beta_i^* : i \in I)$ so that, for each i, $T_i^* = \{t_i^*\}$ and $\text{marg}_{S_{-i}} \beta_i^*(t_i^*) = p_i$ where $p_i \in \Delta(S_{-i})$ so that $p_i \neq \mu_i$. Write $\delta_i^{*,m} : T_i^* \to H_i^m$ for the mapping from the type t_i^* to the m^{th}-order belief and, similarly, write $\delta_i^* : T_i^* \to H_i^\infty$ for the mapping from the type t_i^* to the associated hierarchy of beliefs. Observe that, for each $s_{-i} \in S_{-i}$, $\delta_i^{*,1}(t_i^*)(s_{-i}) = p_i(s_{-i})$. Moreover, for each $m \geq 2$, each $s_{-i} \in S_{-i}$, and each Borel $E_{-i}^{m-1} \subseteq \prod_{n=1}^{m-1} H_{-i}^n$,

$$\delta_i^{*,m}(t_i^*)(\{s_{-i}\} \times E_{-i}^{m-1}) = \begin{cases} p_i(s_{-i}) & \text{if } (\delta_{-i}^{*,1}(t_{-i}^*), \ldots, \delta_{-i}^{*,m-1}(t_{-i}^*)) \in E_{-i}^{m-1} \\ 0 & \text{otherwise.} \end{cases}$$

Write

$$H_i(t_i^*) = \{t_i \in T_i : \delta_i(t_i) = \delta_i^*(t_i^*)\}.$$

for the set of types $t_i \in T_i$ that induce the same hierarchies of beliefs as t_i^*. Similarly, write

$$H_i^m(t_i^*) = \{t_i \in T_i : \delta_i^m(t_i) = \delta_i^{*,m}(t_i^*)\}$$

for the set of types $t_i \in T_i$ that induce the same m^{th}-order beliefs as t_i^*. Of course,

$$H_i(t_i^*) = \bigcap_{m \geq 1} H_i^m(t_i^*).$$

Suppose, contra hypothesis, that $\mathcal{T} = (S_{-i}, T_i, \beta_i : i \in I)$ is hierarchy-complete. Then, for each i, $H_i(t_i^*) \neq \emptyset$. Because $\mathcal{T}$ is a level-k type structure, there exists a map $N_i : H_i(t_i^*) \to \mathbb{N}_+$ so that (i) for each $t_i \in H_i(t_i^*)$, $t_i \in T_i^{N_i(t_i)}$, but (ii) if $N_i(t_i) > m \geq 1$, then $t_i \notin T_i^m$. Observe that $N_i(t_i) \geq 2$, since $t_i \in H_i(t_i^*)$ implies $\delta_i^1(t_i) = p_i \neq \mu_i$. (See Lemma A.1.) Choose a profile of types $(t_i : i \in I) \in \prod_{i \in I} H_i(t_i^*)$ so that, for each player i, there is no $u_i \in H_i(t_i^*)$ with $N_i(u_i) < N_i(t_i)$. Without loss, suppose $i = 1$ is a player with $N := N_1(t_1) \leq N_i(t_i)$ for all $i \in I$. Thus,

$$T_{-1}^{N-1} \cap \bigcap_{m \geq 1} H_{-1}^m(t_{-1}^*) = T_{-1}^{N-1} \cap H_{-1}(t_{-1}^*) = \emptyset.$$

We will use this fact to show that $\delta_1(t_1) \neq \delta_1^*(t_1^*)$, contradicting that $t_1 \in H_1(t_1^*)$.

Suppose, contra hypothesis that, for each m, $\delta_1^m(t_1) = \delta_1^{*,m}(t_1^*)$. Let

$$F_{-1}^{m-1} = \{(\delta_{-1}^{*,1}(t_{-1}^*), \ldots, \delta_{-1}^{*,m-1}(t_{-1}^*))\}$$

and observe that the set is Borel. (Note, F_{-1}^{m-1} is the set of $n = 1, \ldots, (m-1)^{th}$-order beliefs of players $i \in I \backslash \{1\}$ that is induced by $\mathcal{T}^*$.) It follows that, for each $m \geq 2$,

$$\delta_1^m(t_1)(\{s_{-1}\} \times F_{-1}^{m-1}) = p_1(s_{-1}).$$

By construction of the maps δ_1^m and the fact that, for each $m \geq 2$, $\delta_1^m(t_1) = \delta_1^{*,m}(t_1^*)$, it follows that

$$\beta_1(t_1)(\{s_{-1}\} \times H_{-1}^{m-1}(t_{-1}^*)) = p_1(s_1)$$

for each $m \geq 2$. But, by construction, $\beta_1(t_1)(\{s_{-1}\} \times T_{-1}^{N-1}) = p(s_{-1})$. So, for each $m \geq 2$,

$$\beta_1(t_1)\left(\{s_{-1}\} \times (H_{-1}^{m-1}(t_{-1}^*) \cap T_{-1}^{N-1})\right) = p_1(s_{-1}).$$

Observe that, for each $m \geq 2$, $H_{-1}^m(t_{-1}^*) \subseteq H_{-1}^{m-1}(t_{-1}^*)$. (See Lemma A.2.) Thus,

$$\begin{aligned}\beta_1(t_1)\left(\{s_{-1}\} \times \bigcap_{m\geq 2}(H_{-1}^{m-1}(t_{-1}^*) \cap T_{-1}^{N-1})\right) &= \lim_{m\to\infty} \beta_1(t_1)\left(\{s_{-1}\} \times (H_{-1}^{m-1}(t_{-1}^*) \cap T_{-1}^{N-1})\right) \\ &= p_1(s_1).\end{aligned}$$

However,

$$\bigcap_{m\geq 2}(H_{-1}^{m-1}(t_{-1}^*) \cap T_{-1}^{N-1}) = \emptyset$$

and, so, for each s_{-1},

$$\beta_1(t_1)(\{s_{-1}\} \times \bigcap_{m\geq 2}(H_{-1}^{m-1}(t_{-1}^*) \cap T_{-1}^{N-1})) = 0.$$

Thus, for each s_{-1}, $p_1(s_{-1}) = 0$, a contradiction.

A.4 Proof of Proposition 5.3

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 A.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 A.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 A.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 $(x, m) \neq (x', m')$ with $\beta_i^*(x, m) = \beta_i^*(x', m')$, 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 A.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. ■

Using Lemmata A.5-A.7, $\mathcal{T}^* = (T_i^*, \beta_i^* : i \in I)$ is a type structure with Polish type sets. 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 A.8. *The type structure $\mathcal{T}^* = (T_i^*, \beta_i^* : i \in I)$ 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. ■

Remark A.1. Because $\mathcal{T}^*$ is a level-k type structure, it is not hierarchy-complete. It is also not type-complete, i.e., the maps β_i^* are not onto. In particular, there is no type t_i^* with $\text{marg}_{S_{-i}} \beta_i^*(t_i^*) \neq \mu_i$ and $\text{Supp}\, \beta_i^*(t_i^*) = S_{-i} \times T^*_{-i}$.

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 covers $\mathcal{C}_i = \{T_i^m : m = 1, 2, \ldots\}$, for each $i \in I$, that jointly satisfy conditions (i)-(ii) of Definition 5.1. 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 A.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 A.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. Measurability 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 $t_i^* \in T_i^{*,m}$, i.e., with $\delta_i^{*,n}(t_i^*) = \delta_i^n(t_i)$. 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 A.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 A.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 A.2-A.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 A.2-A.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 $\nu_i(S_{-i} \times (T_{-i}^* \backslash T_{-i}^{*,m-1})) = 0$. Since $\beta_i(t_i)(S_{-i} \times T_{-i}^{m-1}) = 1$, this is a well-defined measure in $\Delta(S_{-i} \times T_{-i})$ and, moreover, $\nu_i(S_{-i} \times T_{-i}^{*,m-1}) = 1$. By construction, there exists a type $t_i^* \in T_i^{*,m}$ 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,

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

where the second line uses the fact that $\nu_i(S_{-i} \times T_{-i}^{*,m-1}) = 1$, 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 $\beta_i(t_i)(S_{-i} \times T_{-i}^{m-1}) = 1$. This establishes $\delta_i^{*,n+1}(t_i^*) = \delta_i^{n+1}(t_i)$. ■

Remark A.2. Proposition 5.3 establishes that there exists a complete level-k type structure (for $\boldsymbol{\mu}$) that induces all hierarchies of beliefs that are induced by any countable level-k type structure (for $\boldsymbol{\mu}$). The result can be extended to non-countable level-k type structures, provided they induce a countable number of finite-order beliefs. (A proof is available upon request.) But, the result stops short of establishing the following conjecture:

> Conjecture: There exist a complete level-k type structure (for $\boldsymbol{\mu}$) that induces all hierarchies of beliefs induced by some level-k type structure (for $\boldsymbol{\mu}$).

This remains an open question.

To understand the difficulty, note that Lemma A.9 delivers maps $f_i^{m,n} : T_i^m \to T_i^{*,m}$ that are Borel measurable and preserve n^{th}-order beliefs. The fact that each $f_i^{m,n}$ is Borel measurable is important. For instance, the fact that each $f_{-i}^{m,1}$ is Borel measurable is used to show that, for each $t_i \in T_i^m$, there is some $t_i^* \in T_i^{*,m}$ that with $\delta_i^2(t_i) = \delta_i^{*,2}(t_i^*)$.

Here, the fact that each $f_i^{m,n}$ is measurable follows from the fact that T_i is countable. We do not know if the result can be extended to any level-k type structure or even any level-k type structure with Polish type sets T_i. We now explain the tradeoff.

Define a multifunction $F_i^{m,n} : T_i^m \twoheadrightarrow T_i^{*,m}$, where

$$F_i^{m,n}(t_i) = (\delta_i^{*,n})^{-1}(\{\delta_i^n(t_i)\})$$

for each $t_i \in T_i^m$. Suppose this multifunction is non-empty. It suffices to show that there is a Borel measurable selection of $F_i^{m,n}$ or, equivalently, that there is a Borel uniformization or a Borel section of $\text{graph}(F_i^{m,n})$.

Standard measurable selection theorems and uniformization theorems cannot be applied here, since the maps β_i^* (and so $\delta_i^{*,n}$) are not continuous. There are other constructions of the canonical structure, which would lead to β_i^* to be continuous. However, they raise other issues for selection/uniformization. The central issues is that, under those constructions, the sets $T_i^{*,m}$ are not compact. (More details are available upon request.) We do not know if there is a construction or alternate proof that circumvents these issues.

Appendix B Proofs for Sections 6-7

B.1 Proofs for Section 6

Lemma B.1. *Let $\mathcal{J}_i : S_i \twoheadrightarrow \Delta(S_{-i})$ be a correspondence with*

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

Then $\mathcal{J}_i(s_i)$ is closed-valued. Moreover, if $s_i \in S_i^1$, then $\mathcal{J}_i(s_i)$ 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 $(\nu_i^1, \nu_i^2, \ldots)$ with each $\nu_i^k \in \mathcal{J}_i(s_i)$. Then, for each ν_i^k and each $r_i \in S_i$, $\tilde{\pi}_i(s_i, r_i, \nu_i^k) \geq 0$. If $(\nu_i^1, \nu_i^2, \ldots)$ 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, $\nu_i \in \mathcal{J}_i(s_i)$ and $\mathcal{J}_i(s_i)$ is closed. ■

Lemma B.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 B.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

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

It suffices to show that each $O(s_i)$ is closed. If so, each $\{s_i\} \times \beta_i^{-1}(O(s_i))$ is Borel. (This uses the fact that β_i is measurable.) As a consequence,

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

is the finite union of Borel sets and so Borel.

Observe that

$$O(s_i) = \{\nu_i \in \Delta(S_{-i} \times T_{-i}) : \text{marg}_{S_{-i}} \nu_i = \mu_i\} = (\underline{\text{proj}}_{S_{-i}})^{-1}(\mathcal{J}_i(s_i)).$$

Since $\text{proj}_{S_{-i}}$ is continuous, so is $\underline{\text{proj}}_{S_{-i}} : \Delta(S_{-i} \times T_{-i}) \to \Delta(S_{-i})$. So, by Lemma B.1, $O(s_i)$ is closed.

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

B.2 Proof of Theorem 7.1

In what follows, we fix an anchor $\boldsymbol{\mu}$ and the sets L_i^m are defined relative to this anchor.

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

(i) $\text{proj}_{S_i} \left(R_i^1 \cap (S_i \times T_i^1)\right) = L_i^1$, *and*

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

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

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 $s_i \in \text{proj}_{S_i} \left(R_i^{m+1} \cap (S_i \times T_i^{m+1})\right)$. Then there exists some $t_i \in T_i^{m+1}$ so that $(s_i, t_i) \in R_i^{m+1}$. As such, $s_i \in \mathbb{BR}_i[\text{marg}_{S_{-i}} \beta_i(t_i)]$. Moreover, $\beta_i(t_i)(R_{-i}^m \cap (S_{-i} \times T_{-i}^m)) = 1$. So, by the induction hypothesis, $\text{marg}_{S_{-i}} \beta_i(t_i)(L_{-i}^m) = 1$. As such, $s_i \in L_i^{m+1}$. ∎

Proof of Theorem 7.1. Part (i) is Proposition B.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 5.2. The proof is by induction on m.

The case of $m = 1$ is part (i) of Proposition B.1. So, assume the result holds for m. By part (ii) of Proposition B.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 (Lemma B.3), we have that, for each $n \leq m$, $\hat{\nu}_i(R_{-i}^n) = 1$.

Step 2: Since the type structure is a complete level-k type structure for $\boldsymbol{\mu}$, 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 $n \leq m$, $\beta_i(t_i)(R_{-i}^n) = 1$, $(s_i, t_i) \in R_i^{m+1}$. ■

B.3 Result for Section 7.3

Lemma B.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 $L_i^n \subseteq S_i^1$. 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 C Proofs for Section 8

In the main text, we stated the following: There exists a weakly decreasing function $J : \{0, 1, 2, \ldots\} \to \{0, \ldots, \kappa\}$ so that, for each i and each $m \geq 0$,

$$S_i^m = \{\underline{x}, \ldots, \underline{x} + J(m)\Delta\}.$$

Moreover, the function J satisfies the following criterion: (i) $J(0) = \kappa$; (ii) $J(m+1) < J(m)$ if $J(m) \geq 1$ and $\underline{x} + J(m)\Delta > \frac{\Delta}{2(1-p)}$; and (iii) $J(m+1) = J(m)$ otherwise. This appendix is devoted to prove this characterization.

C.1 Key Results

To prove the claim, it will be useful to introduce standard terminology: Fix some non-empty set $X_{-i} \subseteq S_{-i}$. Say s_i is **dominated given** X_{-i} if there exists a $\sigma_i \in \Delta(S_{-i})$ so that $\pi_i(\sigma_i, s_{-i}) > \pi_i(s_i, s_{-i})$, for each $s_{-i} \in X_{-i}$. Say s_i is **justifiable given** X_{-i} if there exists a ν_i with $s_i \in \mathbb{BR}_i[\nu_i]$ and $\nu_i(X_{-i}) = 1$. A standard result is that s_i is justifiable given X_{-i} if and only if it is not dominated given X_{-i}.

The claim follows from the following three lemmata:

Lemma C.1. *Let $s^* = (s_1^*, \ldots, s_{|I|}^*)$ be such that $s_1^* = \cdots = s_{|I|}^*$ Then, $(s_1^*, \ldots, s_{|I|}^*)$ is a Nash equilibrium if and only if either $s_i^* = \underline{x}$ or $s_i^* = \underline{x} + k\Delta \leq \frac{\Delta}{2(1-p)}$.*

Lemma C.2. *Fix some $\underline{x} + k\Delta > \min\{\underline{x}, \frac{\Delta}{2(1-p)}\}$. Then, $\underline{x} + k\Delta$ is dominated given $\{\underline{x}, \underline{x} + \Delta, \ldots, \underline{x} + k\Delta\}^{|I|-1}$.*

Lemma C.3. *Fix some $\underline{x} + k\Delta > \underline{x}$. If $\underline{x} + k\Delta$ is justifiable given $\{\underline{x}, \underline{x} + \Delta, \ldots, \underline{x} + \ell\Delta\}^{|I|-1}$, then $\underline{x} + (k-1)\Delta$ is justifiable given $\{\underline{x}, \underline{x} + \Delta, \ldots, \underline{x} + \ell\Delta\}^{|I|-1}$.*

C.2 Proof of Key Results

Proof of Lemma C.1. Fix $s_1^* = \cdots = s_{|I|}^* = \underline{x} + k\Delta$. Then s^* is a Nash equilibrium if and only if, for each $i \in I$ and each each $j \in \{0, \ldots, \kappa\}$,

$$\pi_i(s_i^*, s_{-i}^*) = B - ((\underline{x} + k\Delta)^2(1-p)^2) \geq B - ((\underline{x} + j\Delta) - p(\underline{x} + k\Delta))^2 = \pi_i(\underline{x} + j\Delta, s_{-i}^*).$$

Thus, s^* is a Nash equilibrium if and only if, for each $j \in \{0, \ldots, \kappa\}$,

$$((\underline{x} + k\Delta)^2(1-p)^2) \leq ((\underline{x} + j\Delta) - p(\underline{x} + k\Delta))^2$$

or equivalently, for each $j \in \{0, \ldots, \kappa\}$,

$$(\underline{x} + k\Delta)^2(1 - 2p + p^2) \leq (\underline{x} + j\Delta)^2 + p^2(\underline{x} + k\Delta)^2 - 2p(\underline{x} + j\Delta)(\underline{x} + k\Delta)$$

or equivalently, for each $j \in \{0, \ldots, \kappa\}$,

$$(\underline{x} + k\Delta)^2 - (\underline{x} + j\Delta)^2 \leq 2p(\underline{x} + k\Delta)^2 - 2p(\underline{x} + j\Delta)(\underline{x} + k\Delta)$$

or equivalently, for each $j \in \{0, \ldots, \kappa\}$,

$$\Delta(k-j)(k+j) + 2\underline{x}(k-j) \leq 2p(\underline{x} + k\Delta)(k-j). \tag{3}$$

First suppose that there is some $j > k$. Equation (3) holds for $j > k$ if and only if

$$\Delta(k+j) + 2\underline{x} \geq 2p(\underline{x} + k\Delta).$$

Thus, Equation (3) holds for all $j > k$ if and only if it holds for $j = k + 1$, i.e., if and only if,

$$\Delta(2k+1) + 2\underline{x} \geq 2p(\underline{x} + k\Delta)$$

or equivalently if and only if

$$2(\underline{x}+k\Delta)(1-p) \geq -\Delta.$$

This trivially holds since $\underline{x} \geq 0$ and $\Delta > 0$.

Second, suppose that there is some $j < k$. Equation (3) holds for $j < k$ if and only if

$$\Delta(k+j) + 2\underline{x} \leq 2p(\underline{x}+k\Delta).$$

Thus, Equation (3) holds for all $j < k$ if and only if either $k = 0$ or $k \geq 1$ and the condition holds for $j = k-1$. If $k \geq 1$, this requires that

$$\Delta(2k-1) + 2\underline{x} \leq 2p(\underline{x}+k\Delta)$$

or equivalently if and only if

$$2(\underline{x}+k\Delta)(1-p) \leq \Delta$$

or equivalently if and only if

$$(\underline{x}+k\Delta) \leq \frac{\Delta}{2(1-p)},$$

as stated. ■

Proof of Lemma C.2. We will show that, for each $s_{-i} = (s_j : j \neq i) \in \{\underline{x}, \underline{x}+\Delta, \ldots, \underline{x}+k\Delta\}^{|I|-1}$, $\pi_i(\underline{x}+(k-1)\Delta, s_{-i}) > \pi_i(\underline{x}+k\Delta, s_{-i})$. To do so, observe that, for each $s_{-i} = (s_j : j \neq i) \in \{\underline{x}, \underline{x}+\Delta, \ldots, \underline{x}+k\Delta\}^{|I|-1}$,

$$\frac{1}{|I|-1}\sum_{j\neq i} s_j \in [\underline{x}, \underline{x}+k\Delta].$$

Thus, it suffices to show that, for each $\alpha \in [\underline{x}, \underline{x}+k\Delta]$,

$$B - (\underline{x}+(k-1)\Delta - p\alpha)^2 > B - (\underline{x}+k\Delta - p\alpha)^2 . \tag{4}$$

Fix some $\alpha \in [\underline{x}, \underline{x}+k\Delta]$ and note that Equation (4) holds if and only if

$$(\underline{x}+(k-1)\Delta - p\alpha)^2 < (\underline{x}+k\Delta - p\alpha)^2 .$$

or equivalently if and only if

$$(\underline{x}+(k-1)\Delta)^2 - 2p\alpha(\underline{x}+(k-1)\Delta) < (\underline{x}+k\Delta)^2 - 2p\alpha(\underline{x}+k\Delta)$$

or equivalently if and only if

$$(\underline{x}+(k-1)\Delta)^2 - (\underline{x}+k\Delta)^2 < 2p\alpha(\underline{x}+(k-1)\Delta) - 2p\alpha(\underline{x}+k\Delta)$$

or equivalently if and only if

$$(k-1)^2\Delta + 2\underline{x}(k-1) - k^2\Delta - 2\underline{x}k < -2p\alpha$$

or equivalently if and only if

$$\Delta(-2k+1) - 2\underline{x} < -2p\alpha$$

or equivalently if and only if

$$p\alpha + \frac{\Delta}{2} < (\underline{x} + k\Delta).$$

But notice that

$$p\alpha + \frac{\Delta}{2} \leq p(\underline{x} + k\Delta) + \frac{\Delta}{2}.$$

So it suffices to show that

$$(\underline{x} + k\Delta)(1-p) > \frac{\Delta}{2},$$

which holds by the assumption that $\underline{x} + k\Delta > \frac{\Delta}{2(1-p)}$. ■

To prove Lemma C.3, it will be useful to have an auxiliary result. Given $\nu_i \in \Delta(S_{-i})$, write

$$\mathbb{E}_{\nu_i}(s_{-i}) = \frac{1}{|I|-1} \sum_{(s_j : j \in I \backslash \{i\}) \in S_{-i}} \left(\sum_{j \in I \backslash \{i\}} s_j \right) \nu_i((s_j : j \in I \backslash \{i\})),$$

for the expected strategy played under the belief ν_i.

Lemma C.4. *Fix $s_i = \underline{x} + k\Delta$ so that $s_i \in \mathbb{BR}_i[\nu_i]$ for some $\nu_i \in \Delta(S_{-i})$. Then there exist a $y \in [0, \kappa]$ so that*

(i) $\underline{x} + y\Delta = \max\{p\mathbb{E}_{\nu_i}(s_{-i}), \underline{x}\}$, and

(ii) k is either $\lfloor y \rfloor$ or $\lceil y \rceil$.

Moreover, if both $\lfloor y \rfloor, \lceil y \rceil \in \{0, \ldots, \kappa\}$, then $k = \lfloor y \rfloor$ only if $y - \lfloor y \rfloor \leq \lceil y \rceil - y$ and $k = \lceil y \rceil$ only if $\lceil y \rceil - y \leq y - \lfloor y \rfloor$.

Proof. Construct an auxiliary function $f_i : [0, \underline{x} + \kappa\Delta] \times \Delta(S_{-i}) \to \mathbb{R}$ so that

$$f_i(x_i, \nu_i) = \sum_{(s_j : j \in I \backslash \{i\}) \in S_{-i}} \left[B - (x_i - \frac{p}{|I|-1} \sum_{j \in I \backslash \{I\}} s_j)^2 \right] \nu_i((s_j : j \in I \backslash \{i\})).$$

Note, for any $\nu_i \in \Delta(S_{-i})$, $f_i(x_i, \nu_i)$ is strictly increasing at x_i (resp. strictly decreasing) provided $x_i < p\mathbb{E}_{\nu_i}(s_{-i})$ (resp. $x_i > p\mathbb{E}_{\nu_i}(s_{-i})$); it is maximized at $x_i^* = p\mathbb{E}_{\nu_i}(s_{-i})$.

Fix some ν_i. If $p\mathbb{E}_{\nu_i}(s_{-i}) \leq \underline{x}$, take $y = 0$ and note that $\lfloor 0 \rfloor = \lceil 0 \rceil = 0$ from which the result follows. So, for the remainder, assume $p\mathbb{E}_{\nu_i}(s_{-i}) > \underline{x}$.

Note, there exists some $y \in (0, \kappa]$ so that $\underline{x} + y\Delta = p\mathbb{E}_{\nu_i}(s_{-i})$. Since, for each $s_i \in S_i$, $\pi_i(s_i, \nu_i) = f_i(s_i, \nu_i)$, we can conclude:

(i) If $k < \lfloor y \rfloor$, then $\pi_i(\underline{x} + k\Delta, \nu_i) < \pi_i(\underline{x} + \lfloor y \rfloor \Delta, \nu_i)$; and

(ii) if $k > \lceil y \rceil$, then $\pi_i(\underline{x} + k\Delta, \nu_i) < \pi_i(\underline{x} + \lceil y \rceil \Delta, \nu_i)$.

As a consequence, if $\underline{x} + k\Delta \in \mathbb{BR}_i[\nu_i]$ then either $k = \lfloor y \rfloor$ or $k = \lceil y \rceil$.

Finally, observe that, for each $z \in \mathbb{R}$, $f_i(p\mathbb{E}_{\nu_i}(s_{-i}) + z, \nu_i) = f_i(p\mathbb{E}_{\nu_i}(s_{-i}) - z, \nu_i)$. So, $f_i(\underline{x} + \lfloor y \rfloor \Delta) \geq f_i(\underline{x} + \lceil y \rceil \Delta)$ if and only if $y - \lfloor y \rfloor \leq \lceil y \rceil - y$. And, similarly, $f_i(\underline{x} + \lfloor y \rfloor \Delta) \leq f_i(\underline{x} + \lceil y \rceil \Delta)$ if and only if $y - \lfloor y \rfloor \geq \lceil y \rceil - y$. ■

Proof of Lemma C.3. Suppose $s_i = \underline{x} + k\Delta > \underline{x}$ is justifiable given $\{\underline{x}, \underline{x} + \Delta, \ldots, \underline{x} + \ell\Delta\}^{|I|-1}$. Using the definition of justifiability and Lemma C.2, it follows that $\ell > k$ and there exists a $\nu_i^0 \in \Delta(S_{-i})$ that satisfies the following properties:

(i) $s_i \in \mathbb{BR}_i[\nu_i^0]$;

(ii) $\operatorname{Supp} \nu_i^0 \subseteq \{\underline{x}, \underline{x} + \Delta, \ldots, \underline{x} + \ell\Delta\}^{|I|-1}$; and

(iii) there must be a player $j \neq i$ and an $s_j \in \{\underline{x} + (k+1)\Delta, \ldots, \underline{x} + \ell\Delta\}$ with $\operatorname{marg}_{S_j} \nu_i^0(s_j) > 0$.

Note, if $s_i - \Delta \in \mathbb{BR}_i[\nu_i^0]$, then we are done. So, throughout, suppose $s_i - \Delta \notin \mathbb{BR}_i[\nu_i^0]$.

Construct maps $f_j : S_j \to S_j$ so that

$$f_j(\underline{x} + n\Delta) = \begin{cases} \underline{x} + (n-1)\Delta & \text{if } n \geq 1 \\ \underline{x} & \text{if } n = 0. \end{cases}$$

and let $f_{-i} : S_{-i} \to S_{-i}$ be the associated product map. Let ν_i^1 be the image measure of f_{-i} under ν_i^0 and observe that $\operatorname{Supp} \nu_i^1 \subseteq \{\underline{x}, \underline{x} + \Delta, \ldots, \underline{x} + \ell\Delta\}^{|I|-1}$.

It suffices to show that $\mathbb{BR}_i[\nu_i^1] \subseteq \{s_i, s_i - \Delta\}$: If $s_i - \Delta = \underline{x} + (k-1)\Delta$ is a best response under ν_i^1, it is justifiable given $\{\underline{x}, \underline{x} + \Delta, \ldots, \underline{x} + \ell\Delta\}^{|I|-1}$. If not, ν_i^1 satisfies the same three properties that ν_i^0 satisfied (properties (i)-(ii)-(iii)). Thus, we can repeat the argument and construct a ν_i^2 that is the image measure of f_{-i} under ν_i^1. By construction, $\operatorname{Supp} \nu_i^2 \subseteq \{\underline{x}, \underline{x} + \Delta, \ldots, \underline{x} + \ell\Delta\}^{|I|-1}$ and $\mathbb{BR}_i[\nu_i^2] \subseteq \{s_i, s_i - \Delta\}$. And so on. The claim is that there must be an M so that $\{s_i - \Delta\} \subseteq \mathbb{BR}_i[\nu_i^M]$. If not, for each m, $\mathbb{BR}_i[\nu_i^m] = \{s_i\}$. But, there exists m with $\nu_i^m(\underline{x}, \ldots, \underline{x}) = 1$ and only $\underline{x} < s_i$ is a best response to $s_{-i} = (\underline{x}, \ldots, \underline{x})$, a contradiction.

With the above in mind, we focus on showing that $\mathbb{BR}_i[\nu_i^1] \subseteq \{s_i, s_i - \Delta\}$. Observe that $s_i = \underline{x} + k\Delta$ for $k \geq 1$; so, there exists some $y^0 > 0$ with $p\mathbb{E}_{\nu_i^0}(s_{-i}) = \underline{x} + y^0\Delta$. Similarly, there exists some $y^1 \in \mathbb{R}$ so that $p\mathbb{E}_{\nu_i^1}(s_{-i}) = \underline{x} + y^1\Delta$. (Note, $y^1 < 0$ only if $p\mathbb{E}_{\nu_i^0}(s_{-i}) < \underline{x}$.) Since $\mathbb{E}_{\nu_i^1}(s_{-i}) \in [\mathbb{E}_{\nu_i^0}(s_{-i}) - \Delta, \mathbb{E}_{\nu_i^0}(s_{-i}))$, $y^1 \in [y^0 - 1, y^0)$.

Observe that $\mathbb{BR}_i[\nu_i^1] \subseteq \{\underline{x} + \lfloor y^1 \rfloor \Delta, \underline{x} + \lceil y^1 \rceil \Delta\}$. (See Lemma C.4.) Thus, it suffices to show that (i) if $\underline{x} + \lfloor y^1 \rfloor \Delta \in \mathbb{BR}_i[\nu_i^1]$, then $\lfloor y^1 \rfloor \in \{k, k-1\}$, and (ii) if $\underline{x} + \lceil y^1 \rceil \Delta \in \mathbb{BR}_i[\nu_i^1]$, then $\lceil y^1 \rceil \in \{k, k-1\}$.

First, suppose that $k = y^0$. Since $y^1 \in [y^0 - 1, y^0)$, it follows that $\lfloor y^1 \rfloor = k - 1$ and $\lceil y^1 \rceil = k$, as required.

Second, suppose that $k = \lfloor y^0 \rfloor < y^0$. If $(\underline{x} + \lfloor y^1 \rfloor \Delta) \in \mathbb{BR}_i[\nu_i^1]$, then $\lfloor y^1 \rfloor$ is either $\lfloor y^0 \rfloor = k$ or $\lfloor y^0 - 1 \rfloor = k - 1$. (This follows from the fact that $y^1 \in [y^0 - 1, y^0)$.) So suppose $(\underline{x} + \lceil y^1 \rceil \Delta) \in \mathbb{BR}_i[\nu_i^1]$. Observe that

$$(k+1) = \lceil y^0 \rceil \geq \lceil y^1 \rceil \geq \lfloor y^0 - 1 \rfloor = (k-1),$$

where the first equality uses the fact that $k = \lfloor y^0 \rfloor < y^0$. Thus, it suffices to show that $\lceil y^1 \rceil \neq (k+1)$: If $\lceil y^0 \rceil = \lceil y^1 \rceil = (k+1)$ then $k = \lfloor y^0 \rfloor \geq \lfloor y^1 \rfloor$ implies $\lfloor y^1 \rfloor = k$. Since $(\underline{x} + \lceil y^1 \rceil \Delta) \in \mathbb{BR}_i[\nu_i^1]$, $y^1 - \lfloor y^1 \rfloor \geq$

$\lceil y^1 \rceil - y^1$ (Lemma C.4) and so

$$y^0 - k > y^1 - k = y^1 - \lfloor y^1 \rfloor \geq \lceil y^1 \rceil - y^1 = (k+1) - y^1 > (k+1) - y^0$$

and so

$$y^0 - \lfloor y^0 \rfloor > \lceil y^0 \rceil - y^0$$

contradicting the fact that $k = \lfloor y^0 \rfloor$ (Lemma C.4).

Third, suppose that $k = \lceil y^0 \rceil > y^0$. If $(\underline{x} + \lceil y^1 \rceil \Delta) \in \mathbb{BR}_i[\nu_i^1]$, then $\lceil y^1 \rceil$ is either $\lceil y^0 \rceil = k$ or $\lceil y^0 - 1 \rceil = k - 1$. (This follows from the fact that $y^1 \in [y^0 - 1, y^0)$.) So suppose $(\underline{x} + \lfloor y^1 \rfloor \Delta) \in \mathbb{BR}_i[\nu_i^1]$. Since $k = \lceil y^0 \rceil \geq \lceil y^1 \rceil \geq \lceil y^0 - 1 \rceil = k - 1$,

$$k \geq \lfloor y^1 \rfloor \geq (k-2).$$

As such, it suffices to show that $\lfloor y^1 \rfloor \neq (k-2)$. Suppose, contra hypothesis, that $\lfloor y^1 \rfloor = (k-2)$. Using the fact that $\lceil y^1 \rceil \geq \lceil y^0 - 1 \rceil = k - 1$, $\lceil y^1 \rceil = k - 1$. Since $(\underline{x} + \lfloor y^1 \rfloor \Delta) \in \mathbb{BR}_i[\nu_i^1]$, $\lceil y^1 \rceil - y^1 \geq y^1 - \lfloor y^1 \rfloor$ (Lemma C.4),

$$(k-1) - y^1 \geq y^1 - (k-2)$$

or

$$k - \frac{3}{2} \geq y^1.$$

Now, using the fact that $s_i \in \mathbb{BR}_i[\nu_i^0]$ and $s_i - \Delta \notin \mathbb{BR}_i[\nu_i^0]$, $y^0 - \lfloor y^0 \rfloor > \lceil y^0 \rceil - y^0$ (Lemma C.4). Since $k = \lceil y^0 \rceil > y^0$, $\lfloor y^0 \rfloor = (k-1)$ and so

$$y^0 - (k-1) > k - y^0$$

or

$$y^0 > k - \frac{1}{2}.$$

But now observe that

$$k - \frac{3}{2} \geq y^1 \geq y^0 - 1 > k - \frac{3}{2},$$

a contradiction. ■

C.3 Eliminating One Strategy at a Time

Recall, $\tilde{x}(\Delta, p) = \frac{\Delta}{2(1-p)}$. The main text eluded to the following:

Proposition C.1. *Suppose $\underline{x} + \kappa\Delta \leq 3\tilde{x}(\Delta, p)$. Then there exists some M so that $J(m) = J(m-1) - 1$ for all $m < M$ and $J(m) = J(M)$ for all $m \geq M$.*

This proposition will follow from the following Lemma.

Lemma C.5. *For each $i \in I$, set $s_i = \underline{x} + k\Delta$ and $s_i^* = \underline{x} + (k+1)\Delta$. If $(s_1, \ldots, s_{|I|})$ is not a Nash equilibrium, then $s_i = \underline{x} + k\Delta$ is a best response given s_{-i}^* if and only if $s_i = \underline{x} + k\Delta \leq \frac{\Delta(2p+1)}{2(1-p)}$.*

To see why the lemma suffices: Recall $J(0) = \kappa$. If $\underline{x} + \kappa\Delta \leq \tilde{x}(\Delta, p)$, then $J(m) = \kappa$ for all m. So suppose $\underline{x} + \kappa\Delta > \tilde{x}(\Delta, p)$. In this case $J(1) < J(0)$. We will have $J(1) = \kappa - 1$ if either $(\underline{x} + (\kappa-1)\Delta, \ldots, \underline{x} + (\kappa-1)\Delta)$ is a Nash equilibrium or $\underline{x} + (\kappa-1)\Delta$ is a best response under $(\underline{x} + \kappa\Delta, \ldots, \underline{x} + \kappa\Delta)$. The former requires

$\underline{x}+\kappa\Delta \leq \tilde{x}(\Delta,p)+\Delta$ and the latter requires $\underline{x}+\kappa\Delta \leq 3\tilde{x}(\Delta,p)$; since $3\tilde{x}(\Delta,p) > \tilde{x}(\Delta,p)+\Delta$, $J(1)=\kappa-1$ if $\underline{x}+\kappa\Delta \leq 3\tilde{x}(\Delta,p)$. If $J(2)=J(1)$, we are done. If not, $J(2)=J(1)-1$ since $\underline{x}+(\kappa-2)\Delta < \underline{x}+(\kappa-1)\Delta$. And so on.

Proof of Lemma C.5. For each $j \in \{0,\dots,\kappa\}$,

$$\pi_i(\underline{x}+k\Delta, s^*_{-i}) \geq \pi_i(\underline{x}+j\Delta, s^*_{-i})$$

if and only if, for each $j \in \{0,\dots,\kappa\}$,

$$((\underline{x}+k\Delta) - p(\underline{x}+(k+1)\Delta))^2 \leq ((\underline{x}+j\Delta) - p(\underline{x}+(k+1)\Delta))^2$$

or equivalently, for each $j \in \{0,\dots,\kappa\}$,

$$(\underline{x}+k\Delta)^2 - 2p(\underline{x}+(k+1)\Delta)(\underline{x}+k\Delta) \leq (\underline{x}+j\Delta)^2 - 2p(\underline{x}+(k+1)\Delta)(\underline{x}+j\Delta)$$

or equivalently, for each $j \in \{0,\dots,\kappa\}$,

$$\Delta(k+j)(k-j) + 2\underline{x}(k-j) \leq 2p(\underline{x}+(k+1)\Delta)(k-j). \tag{5}$$

First, observe that Equation (5) holds for each $j > k$ if and only if

$$\Delta(k+j) + 2\underline{x} \geq 2p(\underline{x}+(k+1)\Delta).$$

Thus, it will hold for all $j > k$ if and only if it holds for $j = k+1$, i.e., if and only if

$$\Delta(2k+1) + 2\underline{x} \geq 2p(\underline{x}+(k+1)\Delta).$$

or if and only if

$$2(1-p)(\underline{x}+k\Delta) \geq \Delta(2p-1).$$

Since $(s_1,\dots,s_{|I|})$ is not a Nash equilibrium $2(1-p)(\underline{x}+k\Delta) > \Delta$. Given that $\Delta \geq (2p-1)\Delta$, the conclusion holds.

Second, observe that Equation (5) holds for each $j < k$ if and only if

$$\Delta(k+j) + 2\underline{x} \leq 2p(\underline{x}+(k+1)\Delta).$$

Thus, it will hold for all $j < k$ if and only if it holds for $j = k-1$, i.e., if and only if

$$\Delta(2k-1) + 2\underline{x} \leq 2p(\underline{x}+(k+1)\Delta).$$

or if and only if

$$2(1-p)(\underline{x}+k\Delta) \leq (1+2p)\Delta,$$

as stated. ■

Appendix D Proofs for Section 9

D.1 Proof of Proposition 9.1

The proof is analogous to Example 9.1: 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 $h^m_{i,\Box}$ so that $h^1_{i,\Box}(\Box_{-i}) = 1$ and $h^{m+1}_{i,\Box}(\Box_{-i}, \ldots, h^m_{-i,\Box}) = 1$. Set $h_{i,\Box} = (h^1_{i,\Box}, h^2_{i,\Box}, \ldots)$. Likewise, for each player i, inductively define h^m_i as follows: First, set $h^1_i = \mu_i$. Second, $h^m_i(\Box_{-i}, h^1_{-i,\Box}, \ldots, h^m_{-i,\Box}) = p \in (0, \mu_i(\Box_{-i}))$. (Note, p does not depend on m.) Set $h_i = (h^1_i, h^2_i, \ldots)$. Proposition 9.1 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 $m \geq 1$,

$$h^{m+1}_i(\Box_{-i}, h^1_{-i,\Box}, \ldots, h^m_{-i,\Box}) = p$$

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

$$E^{m+1}_i := (\rho^{m+1}_i)^{-1}(\{\Box_{-i}, h^1_{-i,\Box}, \ldots, h^m_{-i,\Box}\}).$$

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

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

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

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

i.e., there exists some type $t_{-i} \in T_{-i}$ with $\delta_{-i}(t_{-i}) = (h^1_{-i,\Box}, h^2_{-i,\Box}, \ldots)$, 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 $t_{i,\Box} \in T_i$ with $\delta_i(t_{i,\Box}) = h_{i,\Box}$.*

Proof. For each $i \in I$, let $\mathcal{C}_i = \{T^m_i : m = 1, 2, \ldots\}$ be a Borel cover so that $(\mathcal{C}_i : i \in I)$ jointly satisfy conditions (i)-(ii) of Definition 5.1. We will show that, for each $m \geq 1$, and each $t_i \in T^m_i$, $\delta^m_i(t_i) \neq h^m_{i,\Box}$. The proof is by induction on m.

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

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

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 m-type and an n-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, $\text{marg}_{S_{-i}}\beta_i(t_i) = \mu_i$ if and only if $t_i \in \{1, 3\}$. Second, $\text{Supp}\,\text{marg}_{T_{-i}}\beta_i(1) = T_{-i}$. 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 $\{T_i^m : m = 1, 2, \ldots\}$ so that $T_i^1 = \{1, 3\}$ and, for each $m \geq 2$, $T_i^m = \{m\}$. This cover is non-partitional. However, any cover must be non-partitional. To see this, fix a cover $\{U_i^m : m = 1, 2, \ldots\}$. Since $\text{Supp}\,\text{marg}_{T_{-i}}\beta_i(1) = T_{-i}$, 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 $U_i^1 \cap U_i^3 \neq \emptyset$. □

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 m-type for one associated cover and an n-type for another associated cover, despite the fact that $m \neq n$.

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, $\text{marg}_{S_{-i}}\beta_i(t_i) = \mu_i$ if and only if $t_i \in \{1, 3\}$. Second, $\text{Supp}\,\text{marg}_{T_{-i}}\beta_i(1) = T_{-i}$. Third, for each $m \geq 1$, $\beta_i(m+1)(S_{-i} \times \{m\}) = 1$.

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

D.3 Finite-Order Belief Type Structures

We begin by introducing a type structure that can capture the idea that a player thinks other players do not have an m^{th}-order belief.

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, for each $i \in I$,

(i) $\tilde{T}_i$ is a metrizable set of types for i,

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

(iii) $\tilde{T}_i \backslash (\tilde{\beta}_i)^{-1}(\{d\}) \neq \emptyset$.

To better understand, observe that, now, there can be a type $\tilde{t}_i \in \tilde{T}_i$ with $\tilde{\beta}_i(\tilde{t}_i) = d$. This type is a "dummy type," that is not associated with a belief. Loosely, think of this type as one that does not "reason." So, if $\beta_i(t_i)$ assigns positive probability to a "dummy type" of player j, then t_i assigns positive probability to the event that player j "does not reason." The fact that each $\tilde{T}_i \backslash (\tilde{\beta}_i)^{-1}(\{d\}) \neq \emptyset$ implies that no player only has "dummy types." We refer to a finitary S-based type structure as, a finitary type structure

Say $(s_i, \tilde{t}_i)$ is **rational** if $\tilde{\beta}_i(\tilde{t}_i) \neq \{d\}$ and satisfies the condition in Definition 6.1. Say $\tilde{t}_i$ **believes** an event E_{-i} if $\tilde{\beta}_i(\tilde{t}_i) \neq \{d\}$ and $\tilde{t}_i$ satisfies the condition in Definition 6.2. We define RmBR analogously to Definition 6.3. 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 $\text{marg}_{S_{-i}} \tilde{\beta}_i(\tilde{t}_i) = \nu_i$.

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

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 $s_i \in S_i$, $s_i \notin \mathbb{BR}_i[f_i(s_i)]$. As a consequence, each strategy can be inconsistent with rationality. This is important: If $\tilde{t}_i$ is a "dummy type," each element of $S_i \times \{\tilde{t}_i\}$ is irrational. We would like a type t_i that assigns positive probability to $(s_i, \tilde{t}_i)$ to retain its marginal belief about behavior, while continuing to assign positive probability to irrationality.

Proof of Proposition D.1. Fix a game with no weakly dominant strategy and an associated finitary type structure that is first-order complete, viz. $\tilde{\mathcal{T}}$. Let $T_i = \tilde{T}_i \backslash (\tilde{\beta}_i)^{-1}(\{d\})$. Since $\tilde{\beta}_i$ is measurable, T_i is a Borel subset of $\tilde{T}_i$. Endow T_i with the relative topology and note that it is metrizable.

To construct the belief maps, it will be useful to first define auxiliary maps. Recall, since the game has 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 $\tilde{\mathcal{T}}$ is first-order complete, there are mappings $\tilde{\tau}_i : S_i \to \tilde{T}_i$ such that, for each s_i, $\tilde{\beta}_i(\tilde{\tau}_i(s_i)) \neq d$ and $\text{marg}_{S_{-i}} \tilde{\beta}_i(\tilde{\tau}_i(s_i)) = f_i(s_i)$. Now, define maps $\tau_i : S_i \times \tilde{T}_i \to S_i \times T_i$ so that

$$\tau_i(s_i, t_i) = \begin{cases} (s_i, t_i) & \text{if } t_i \in T_i \\ (s_i, \tilde{\tau}_i(s_i)) & \text{if } t_i \notin T_i. \end{cases}$$

Observe that

$$\tau_i(\{s_i\} \times (\tilde{T}_i \backslash T_i)) = \begin{cases} \emptyset & \text{if } \tilde{T}_i \backslash T_i = \emptyset \\ \{(s_i, \tilde{\tau}_i(s_i))\} & \text{if } \tilde{T}_i \backslash T_i \neq \emptyset. \end{cases}$$

As a consequence, τ_i is measurable: For each Borel $E_i \subseteq S_i \times T_i$,

$$(\tau_i)^{-1}(E_i) = \left(\bigcup_{s_i:(s_i,\tilde{\tau}_i(s_i))\in E_i \text{ and } \tilde{\tau}_i(s_i)\notin T_i} (\{s_i\} \times (\tilde{T}_i \backslash T_i)) \right) \cup E_i$$

is the finite union of Borel sets and so Borel.

Let $\beta_i : T_i \to \Delta(S_{-i} \times T_{-i})$ so that, for each $t_i \in T_i$, $\beta_i(t_i)$ is the image measure of $\tilde{\beta}_i(t_i)$ under τ_{-i}. Note, β_i is measurable, since $\beta_i = \underline{\tau}_{-i} \circ \tilde{\beta}_i$.

To better understand the construction, consider the set of types

$$T_i^\circ = \{t_i \in T_i : \tilde{\beta}_i(t_i)(S_{-i} \times T_{-i}) = 1\}.$$

Note, for each $t_i \in T_i^\circ$ and each Borel $E_{-i} \subseteq S_{-i} \times T_i$, $\beta_i(t_i)(E_{-i}) = \tilde{\beta}_i(t_i)(E_{-i})$. Consider instead some $t_i \in T_i \backslash T_i^\circ$. Then, there exists some $j \neq i$ so that $\text{marg}_{S_{-j}\times T_{-j}} \tilde{\beta}_i(t_i)(S_j \times (\tilde{T}_j \backslash T_j)) > 0$. As such, there will be $(s_j, \tau_j(s_j))$ so that $\text{marg}_{S_j \times T_j} \beta_i(t_i)(\{(s_j, \tilde{\tau}_j(s_j))\}) > 0$. The key is that each $(s_j, \tilde{\tau}_j(s_j))$ is irrational.

With this in mind, we complete the proof by showing the following: For each $m \geq 1$ and each $t_i \in T_i$,

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

(ii) if $t_i \in T_i \backslash T_i^\circ$, $(S_i \times \{t_i\}) \cap R_i^{m+1} = (S_i \times \{t_i\}) \cap \tilde{R}_i^{m+1} = \emptyset$.

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 $(s_i, t_i) \in \tilde{R}_i^1$.

Now suppose that $t_i \in T_i \backslash T_i^\circ$. Since $\tilde{R}_{-i}^1 \subseteq S_{-i} \times T_{-i}$ and $\tilde{\beta}_i(t_i)(S_{-i} \times T_{-i}) < 1$, $(S_i \times \{t_i\}) \cap \tilde{R}_i^2 = \emptyset$. Similarly, there must be some $j \neq i$ and some $s_j \in S_j$ so that $\text{marg}_{S_j \times T_j} \beta_i(t_i)(\{(s_j, \tilde{\tau}_j(s_j))\}) > 0$. Since $(s_j, \tilde{\tau}_j(s_j)) \notin R_j^1$, it follows that $(S_i \times \{t_i\}) \cap R_i^2 = \emptyset$.

$m \geq 2$: Suppose the claim holds for $m \geq 1$. Fix a type $t_i \in T_i$. If $t_i \notin T_i^\circ$ then the claim follows from the fact that $(S_i \times \{t_i\}) \cap R_i^2 = (S_i \times \{t_i\}) \cap \tilde{R}_i^2 = \emptyset$. So suppose $t_i \in T_i^\circ$. Note $\beta_i(t_i)$ believes R_{-i}^m if and only if $\tilde{\beta}_i(t_i)$ believes R_{-i}^m. (This is by construction.) By the induction hypothesis, $R_{-i}^m = \tilde{R}_{-i}^m$, as required. ■

D.4 Fine Grid

Finite Game Consider the beauty contest. For any given $m \geq 1$, there exists a ε-fine grid so that the m-rationalizable strategies are strictly contained in the $(m-1)$-rationalizable strategies. To see this fix a sequence $((\Delta_n, \kappa_n) : n \geq 1)$ where, for each n, $\Delta_n \kappa_n = \overline{x}$ and $\lim_{n\to\infty} \Delta_n = 0$. For a given grid (Δ_n, κ_n), we eliminate a strategy on round m if $\kappa_n \geq m$ and $\underline{x} + (\kappa_n - m + 1)\Delta_n > \Delta_n \frac{|I|-2p}{2|I|(1-p)}$. For each m, there exists some $\underline{N}(m)$ so that $\kappa_n \geq m$ for all $n \geq \underline{N}(m)$. (This follows from the fact that $\kappa_n \Delta_n = \overline{x} - \underline{x} > 0$ and $\lim_{n\to\infty} \Delta_n = 0$.) Now observe that there is an $\overline{N}(m) \geq \underline{N}(m)$ so that, for each $n \geq \overline{N}(m)$,

$$\overline{x} > \Delta_n \frac{|I| - 2p}{2|I|(1-p)} + \Delta_n(m-1).$$

So, for any $n \geq \overline{N}(m)$, the m-rationalizable strategies are a strict subset of the $(m-1)$-rationalizable strategies in the game associated with the grid (Δ_n, κ_n).

Infinite Game Suppose instead that, in the beauty contest, $S_i = [\underline{x}, \overline{x}]$ where $\overline{x} > \underline{x} \geq 0$. We will argue that, for each $m \geq 1$, $[\underline{x}, \overline{x}) \subseteq S_i^m$.

Fix some $s_i \in [\underline{x}, \overline{x})$ with $s_i \neq 0$. Observe that there exists some

$$\alpha \in \left(1, \min\left\{\frac{|I| - p}{p|I| - p}, \frac{\overline{x}}{s_i}\right\}\right).$$

Note, since $\alpha < \frac{\overline{x}}{s_i}$, $\alpha s_i \in [\underline{x}, \overline{x})$. Suppose ν_i assigns probability 1 to $(\alpha s_i, \ldots, \alpha s_i)$. To show that $s_i \in \mathbb{BR}_i[\nu_i]$ it suffices to show that $\pi_i(s_i, \nu_i) = 1$. Notice that $\alpha s_i > s_i$ since $\alpha > 1$ and $s_i > 0$. So, if the target is smaller than s_i, $\pi_i(s_i, \nu_i) = 1$. In fact, the target is

$$p \frac{s_i + (|I| - 1)\alpha s_i}{|I|}.$$

Since $\alpha > \frac{|I|-p}{p|I|-p}$, s_i is strictly higher than the target, as desired.

Next, suppose that $s_i = 0$. Then $s_i = \underline{x} = 0$. In that case, let ν_i assigns probability 1 to $(\underline{x}, \ldots, \underline{x})$ so that $\pi_i(s_i, \nu_i) > 0$. If i instead chooses $y_i > \underline{x} = 0$, $\pi_i(y_i, \nu_i) = 0$: This follows since $y - \frac{py}{|I|} > \frac{py}{|I|}$. (Here we use the fact that $|I| \geq 3$.)

If follows that $[\underline{x}, \overline{x}) \subseteq S_i^1$. Moreover, we showed that each $s_i \in [\underline{x}, \overline{x})$ is a best response under a ν_i with $\nu_i([\underline{x}, \overline{x})^{|I|-1}) = 1$. Since $[\underline{x}, \overline{x})^{|I|-1} \subseteq S_{-i}^1$, the claim follows by induction.

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.

Arad, Ayala and Ariel Rubinstein. 2012. "The 11-20 money request game: A level-k reasoning study." *American Economic Review* 102(7):3561–73.

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.

Breitmoser, Yves. 2012. "Strategic Reasoning in p-Beauty Contests." *Games and Economic Behavior* 75(2):555–569.

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.

Burchardi, Konrad B and Stefan P Penczynski. 2014. "Out of your mind: Eliciting individual reasoning in one shot games." *Games and Economic Behavior* 84:39–57.

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. 2013. "Robust Multiplicity with a Grain of Naivite." http://tuvalu.santafe.edu/~willemien.kets/.

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

Kneeland, Terri. 2015. "Identifying higher-order rationality." *Econometrica* 83(5):2065–2079.

Ledoux, Alain. 1981. "Concours résultats complets." *Les victimes se sont plu à jouer le* 14:10–11.

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.

Seel, Christian and Elias Tsakas. 2017. "Rationalizability and Nash equilibria in Guessing Games." *Games and Economic Behavior* 106:75–88.

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.