Two Approaches to Iterated Reasoning in Games

Adam Brandenburger* Amanda Friedenberg† Terri Kneeland‡

December 9, 2020

Abstract

Level-k analysis and epistemic game theory are two different ways of investigating iterative reasoning in games. This paper explores the relationship between these two approaches. An important difference between them is that level-k analysis begins with an exogenous anchor on the players' beliefs, while epistemic analysis begins with arbitrary epistemic types (hierarchies of beliefs). To close the gap, we develop the concept of a level-k epistemic type structure, that incorporates the exogenous anchor. We also define a complete level-k type structure where the exogenous anchor is the only restriction on hierarchies of beliefs. One might conjecture that, in a complete structure, the strategies that can be played under rationality and $(m-1)^{th}$-order belief of rationality are precisely those strategies played by a level-k player, for any $k \geq m$. In fact, we prove that the strategies that can be played are the m-rationalizable strategies (i.e., the strategies that survive m rounds of elimination of strongly dominated strategies). This surprising result says that level-k analysis and epistemic game theory are two genuinely different approaches, with different implications for inferring the players' reasoning about rationality from their observed behavior.

1 Introduction

Two separate literatures have grown up around iterative reasoning in games: the level-k literature and epistemic game theory. From the outset, the two literatures have had different priorities. The level-k literature has focused on experiments. (See, e.g., Stahl and Wilson, 1994, 1995, Nagel, 1995, Costa-Gomes, Crawford and Broseta, 2001, and Camerer, Ho and Chong, 2004.) It seeks to provide a simple model of iterative reasoning that best describes the data. The epistemic game theory literature has focused on solution concepts. It seeks to seeks to identify assumptions on iterative reasoning that yield important solution concepts—e.g., rationalizability, Nash equilibrium, or correlated equilibrium.[1] (See e.g., Bernheim, 1984, Pearce, 1984, Aumann, 1987, Brandenburger and Dekel, 1987, and Tan and Werlang, 1988.)

Despite the different priorities, the two literatures have an obvious similarity in terms of their interest in iterated reasoning. The aim of this paper is to narrow the gap between how these two literatures understand the concept of "iterated reasoning." We begin by reviewing the architecture of the two literatures with an eye toward understanding the relationship.

*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 Arizona, amanda@amandafriedenberg.org, amandafriedenberg.org

‡Department of Economics, University College London, t.kneeland@ucl.ac.uk, www.tkneeland.com

[1] A more recent strand of the literature seeks to identify new solution concepts from models of iterated reasoning.

The level-k literature provides an implicit model of iterated reasoning. It labels a particular player with a parameter that describes the number of levels of reasoning in which that player engages. To do so, it begins with an exogenous distribution about how the game is played. This is referred to as an "anchor." (The anchor is often—but not always—assumed to be uniform.) A level-0 (L0) player is assumed to play according to the anchor.[2] A level-1 (L1) player assumes that all other players are L0 and chooses a strategy that is expected-payoff maximizing given the level-0 play (i.e., the anchor). A level-2 (L2) player assumes that all other players are L1 and chooses a strategy that is expected-payoff maximizing under some probability distribution on L1 strategies. And so on for higher-level players. Note, the choice of anchor determines what is (or isn't) level-k behavior.

The epistemic literature provides an explicit description of iterated reasoning. This involves two steps: first, describing the players' hierarchies of beliefs about the play of the game and, second, imposing so-called epistemic conditions. The epistemic conditions formalize the players' degrees of "strategic sophistication." But, the formalization first requires specifying the players' hierarchies of beliefs.

As an illustration of the epistemic approach, suppose that there are two players Ann and Bob. Ann is assumed to have a first-order belief, which now can be any probability distribution on the strategies of Bob. Ann also has a second-order belief, which is a joint probability distribution on Bob's strategies and Bob's first-order belief. And so on for higher-order levels of beliefs. A player's epistemic type is the player's entire hierarchy of beliefs (about the strategies played). Note, in principle, the players can hold any hierarchies of beliefs whatsoever. The epistemic conditions serve two roles. First, they link the hierarchies of beliefs to behavior. For instance, consider the epistemic condition that a player is rational, e.g., Ann chooses a strategy that is expected-payoff maximizing given her belief about how Bob plays. This restricts the behavior of Ann that is associated with a given epistemic type. Second, the epistemic conditions restrict the players' hierarchies of beliefs. For instance, consider the condition of "rationality and first-order belief of rationality." This requires that Ann assign probability 1 to Bob's rationality—i.e., to the rational strategy-type pairs of Bob. This is a restriction on Ann's second-order belief, i.e., her belief about the strategies and first-order belief of Bob.

Note that the level-k analysis begins with an exogenous anchor on the beliefs, whereas, a priori, the epistemic analysis allows for an arbitrary space of hierarchies of beliefs. So, in order to bridge the gap between the level-k and epistemic analyses, we need to represent the exogenous anchor within an epistemic model. This is done by exogenously restricting the hierarchies of beliefs that the players consider possible (i.e., prior to any restrictions due to epistemic conditions). Formally, we restrict attention to what we call "level-k type structures." These are epistemic type structures in which the epistemic types can be decomposed into 1-types, 2-types, and so on. The 1-types are epistemic types whose first-order beliefs are, by assumption, constrained to match the anchor. The 2-types assign probability 1 to 1-types. The 3-types assign probability 1 to 2-types. And so on. Of special interest is a "complete level-k type structure," where the only restrictions on beliefs comes from the anchor.[3]

What are the behavioral implications of assuming rationality and $(m-1)$th-order belief of rationality in a (complete) level-k type structure? How do these implications relate to the typical level-k solution concept? One might conjecture that, in a complete level-k type structure, the strategies that can be played

[2]There are two interpretations of this L0 player. One is that the player does not reason at all, but chooses a mixed strategy that corresponds to the anchor. The second is that the player does not exist but, instead, serves as a way to anchor the beliefs of other players.

[3]Formally, any belief consistent with the anchor is represented by an epistemic type in the complete level-k type structure. Incomplete level-k type structures may involve other exogenous restrictions on the players' hierarchies of beliefs.

under rationality and $(m-1)$th-order belief of rationality are precisely those strategies played by an Lk player, for some $k \geq m$. Somewhat surprisingly, this is not true. Instead, Theorem 7.1 shows: In a complete level-k type structure, the strategies consistent with rationality and $(m-1)$th-order belief of rationality are the m-rationalizable strategies.[4] So, while the anchor is important from the perspective of level-k behavior, it is immaterial from the perspective of RmBR in a complete level-k type structure.

To put this result in context, it will be useful to contrast it with a standard result in epistemic game theory. In a complete type structure, the strategies consistent with rationality and $(m-1)$th-order belief of rationality are the m-rationalizable strategies. (See Battigalli, Friedenberg and Siniscalchi, 2012, Chapter 6.) A complete type structure is one that induces all possible beliefs. (There is no anchoring restriction.) So, our result shows that the exogenous restriction on the hierarchies of beliefs imposed by a complete level-k type structure does not do any work, despite the fact that all hierarchies are anchored by level-0 behavior. Rationality and $(m-1)$th-order belief of rationality gives no sharper prediction than the m-rationalizable strategies.

That said, if we focus exclusively on k-types within a complete level-k type space, then the strategies that survive rationality and $(k-1)^{\text{th}}$-order belief of rationality are exactly the Lk strategies. (See Theorem 6.1.) This result provides a way to understand the level-k solution concept: The concept is best understood by changing the epistemic conditions across the strategy-type pairs. We can understand the solution concept by associating 1-types with the epistemic condition of rationality, 2-types with the epistemic condition of rationality and first-order belief of rationality, and so on. But, there is an important caveat to this interpretation. In a complete type structure, we must have types that are both k-types and ℓ-types for $k \neq \ell$. (Formally, in a complete type structure, the decomposition into 1-types, 2-types, etc., is not a partition but a cover.) As such, there is no unique assignment of types to epistemic conditions.

These results highlight that the level-k solution concept and epistemic game theory are two genuinely different approaches. First, while level-k behavior is sensitive to the anchor, RmBR is not. Second, epistemic game theory specifies epistemic conditions independent of the hierarchies of beliefs that the players hold, while the level-k analysis effectively allows the epistemic conditions to depend on the hierarchies of beliefs.

A comparison of level-k analysis and epistemic game theory also offers the possibility of connecting observable behavior (per level-k analysis) with unobservable epistemic states—that is, with levels of reasoning about rationality. In particular, if we know the relationship between level-k behavior and levels of reasoning about rationality (in the epistemic sense), we can observe behavior and deduce the associated level of rationality.

In the context of a complete type structure—where no anchoring restriction is applied—the relationship between behavior and the associated level of rationality is already known. In this case, the best we can do is identify a player's maximum level of reasoning about rationality: If a player chooses an action that is consistent with rationality and $(m-1)$th-order belief of rationality but not consistent with mth-order belief of rationality then, we can identify the player's maximum level of reasoning about rationality as m. As such, if a player chooses an action that is m-rationalizable but not $(m+1)$-rationalizable, we can identify the player's maximum level of reasoning about rationality as m.

One might conjecture that we can achieve tighter identification by instead looking at a complete level-k type structure. However, in such a type structure, rationality and $(k-1)$th order belief of rationality does not characterize the level-k solution set, but instead gives rise to the k-rationalizable strategies. As such,

[4] The m-rationalizable strategies are those that survive m rounds of elimination of strongly dominated strategies.

the implications of rationality and $(k-1)$th order belief of rationality are exactly the same irrespective of whether the players' have a complete type structure vs. a complete level-k type structure. As such, observing a level-k strategy is informative of a player's maximum level of reasoning about rationality being k only if it is a k-rationalizable but not $(k+1)$-rationalizable strategy. This is not a characteristic of two popular games used in the level-k literature: the 11-20 game (Arad and Rubinstein, 2012) and the Beauty Contest game (Nagel, 1995). Thus, knowing the level-k distribution from behavior in games such as these will not be informative of the distribution of levels of reasoning about rationality.

2 The Game and Solution Concepts

We begin with mathematical preliminaries used throughout the paper. Given a metrizable set Ω, endow Ω with the Borel σ-algebra. The set of Borel probability measures on Ω is $\Delta(\Omega)$; endow $\Delta(\Omega)$ with the topology of weak convergence. Given an index set I and a collection of metrizable sets $(\Omega_i : i \in I)$, write $\Omega_{-i} = \prod_{j \in I \setminus \{i\}} \Omega_j$ and $\Omega = \prod_{i \in I} \Omega_j$. Endow the product of metrizable spaces with the product topology.

Throughout the paper, we fix a finite game $G = (S_i, \pi_i : i \in I)$, where I is the set of players, S_i is player i's strategy set, and $\pi_i : \prod_{j \in I} S_j \to \mathbb{R}$ is player i's payoff function. Each player has at least two strategies, i.e., $|S_i| \geq 2$. Extend $\pi_i : S_i \times \Delta(S_{-i}) \to \mathbb{R}$ in the usual way. Given some $\mu_i \in \Delta(S_{-i})$, write

$$\mathbb{BR}_i[\mu_i] := \{s_i \in S_i : \pi_i(s_i, \mu_i) \geq \pi_i(r_i, \mu_i), \text{ for each } r_i \in S_i\}.$$

So, $\mathbb{BR}_i[\mu_i]$ is the set of strategies of player i that are optimal under μ_i.

We will be interested in two iterative solution concepts: rationalizability and level-k. We begin with rationalizability. Set $S_i^0 = S_i$. Assume the sets S_i^m have been defined. Write $s_i \in S_i^{m+1}$ if and only if there exists some $\nu_i \in \Delta(S_{-i})$ with: (i) $s_i \in \mathbb{BR}_i[\nu_i]$, and (ii) $\nu_i(S_{-i}^m) = 1$. Set $S_i^\infty = \bigcap_{k \geq 0} S_i^k$. The set S_i^m (S_i^∞) is the set of **m-rationalizable strategies** (resp. **rationalizable strategies**).

The level-k solution concept begins by exogenously specifying an anchor, $\mu = (\mu_i : i \in I) \in \prod_{i \in I} \Delta(S_{-i})$. Then inductively define sets $L_i^k[\mu]$ as follows: Set $L_i^1[\mu] = \mathbb{BR}_i[\mu_i]$. Assume the sets $L_i^k[\mu]$ have been defined. Let $L_i^{k+1}[\mu]$ be the set of strategies s_i so that there exists some $\nu_i \in \Delta(S_{-i})$ with: (i) $s_i \in \mathbb{BR}_i[\nu_i]$, and (ii) $\nu_i(L_{-i}^k[\mu]) = 1$. Call the set $L^k[\mu]$ the **level-k behavior for μ**. Call the profile $(L^1[\mu], L^2[\mu], \ldots)$ the **level-k solution concept for μ**. When it is clear which μ is being referenced, we often drop the reference to μ, writing only L^k or referring to L^k as level-k behavior.

To better understand the relationship between the two concepts, it will be illustrative to consider an example. Figure 2.1 depicts an Undercutting Game. Each player chooses an integer between 1 and $X \geq 4$. If the other player chooses a number $s_{-i} \geq 2$, player i has an incentive to undercut her opponent by exactly 1. If the other player chooses the lowest number $s_{-i} = 1$, then i has an incentive to match her opponent.

Begin with the solution concept of rationalizability. Since each player has an incentive to undercut the other player by 1, each $s_i \leq X - 1$ is a best response to the strategy $s_{-i} = s_i + 1$. On the other hand, the highest strategy $s_i = X$ is strongly dominated by a mixed strategy; as such, it is not a best response under any pure strategy. Thus, the 1-rationalizable strategy set is $\{1, 2, 3, \ldots, X-1\}$. Inductively,

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

	1	2	3	$\cdots$	$X-1$	X
1	1,1	1,0	0,0	$\cdots$	0,0	0,0
2	0,1	0,0	1,0	$\cdots$	0,0	0,0
3	0,0	0,1	0,0	$\cdots$	0,0	0,0
4	0,0	0,0	0,1	$\cdots$	0,0	0,0
$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$
$X-2$	0,0	0,0	0,0	$\cdots$	1,0	0,0
$X-1$	0,0	0,0	0,0	$\cdots$	0,0	1,0
X	0,0	0,0	0,0	$\cdots$	0,1	0,0

Figure 2.1: Undercutting Game

The level-k solution concept begins with an anchor $\mu = (\mu_1, \mu_2)$. For now, choose the anchor so that each player assigns probability 1 to the largest integer less than or equal to the arithemetic mean, i.e., $\mathbb{E}(x) = \lfloor \frac{1}{X} \sum_{x=1}^{X} x \rfloor$. Thus $L_i^1[\mu] = \{\mathbb{E}(x) - 1\}$. In fact, inductively,

$$L_i^k[\mu] = \begin{cases} \{\mathbb{E}(x) - k\} & \text{if } \mathbb{E}(x) - 1 \geq k \\ \{1\} & \text{if } k \geq \mathbb{E}(x). \end{cases}$$

Notice that, for each m, $L_i^m[\mu] \subseteq S_i^m$. This is true more generally, i.e., beyond the current example. (See, e.g., Crawford, Costa-Gomes and Iriberri (2013).) Since the sets S_i^m are (weakly) shrinking, this implies that, for each m,

$$\bigcup_{k \geq m} L_i^m[\mu] \subseteq S_i^m. \tag{1}$$

However, the inclusion may be strict. To see this, observe that, for each $m \leq \mathbb{E}(x) - 1$,

$$\bigcup_{k \geq m} L_i^m[\mu] = \{1, \ldots, \mathbb{E}(x) - m\} \subsetneq \{1, 2, 3, \ldots, X - m\} = S_i^m.$$

Thus, the inclusion is strict, whenever $m \leq X - 2$.

3 Epistemic Games and Epistemic Conditions

Much as in the level-k approach, the epistemic approach begins by specifying the players' hierarchies of beliefs about how the game is played. Unlike the level-k approach, it does so by expanding the description of the strategic situation to include the players' hierarchies of beliefs. Those hierarchies are described by a type structure, in the sense of Harsanyi (1967).

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.

So, an S-based type structure has two ingredients. First, each player has a set of types. As in the level-k approach, those types do not impact the payoff functions. Instead, they describe the hierarchies of beliefs that players hold about the play of the game. The epistemic approach specifies these hierarchies formally; this is done via a belief map. The belief map associates each type t_i with a belief $\beta_i(t_i)$ on the strategy-type pairs of other players. By doing so, each type induces a hierarchy of beliefs about the play of the game: Type t_i's first-order belief is a belief about the strategies of other players, i.e., it is the belief $\text{marg}_{S_{-i}}\beta_i(t_i)$. But, because each $\beta_i(t_i)$ is a belief about the strategy-type pairs, and the types of other players induce first-order beliefs, type t_i induces a second-order belief about the pairs of strategies and first-order beliefs of the other players. And so on.[5] The following example illustrates this fact.

Example 3.1. Refer back to the undercutting game and take $X = 3$. Consider an S-based type structure with type sets $T_i = \{t_i^1, t_i^2, t_i^3\}$ and belief maps defined as follows:

$$\beta_i(t_i^1)(\{(1, t_{-i}^1)\}) = 1 \qquad \beta_i(t_i^2)(\{(2, t_{-i}^3)\}) = 1 \qquad \beta_i(t_i^3)(\{(3, t_{-i}^3)\}) = 1.$$

Each type t_i^k has a first-order belief that assigns probability 1 to the other player choosing the strategy $s_{-i} = k$. With this, type t_i^2 has a second-order belief that assigns probability 1 to "the other player chooses 2 and believes I choose 3," while type $t_i^k \in \{t_i^1, t_i^3\}$ has a second-order belief that assigns probability 1 to "the other player chooses k and believes I choose k." And so on, inductively. □

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

Definition 3.2. Call an S-based type structure $\mathcal{T} = (S_{-i}, T_i, \beta_i : i \in I)$ **complete** if, for each i, β_i is onto.

The canonical constructions of a so-called universal type structure (e.g,. Mertens and Zamir, 1985 or Brandenburger and Dekel, 1993) are each complete. As such, there exists a complete type structure. But there are also complete type structures that are distinct from these constructions.

We think of an **epistemic game** as a pair $(G, \mathcal{T})$. It describes the play of the game and the players' hierarchies of beliefs. Because we have fixed the game G, we will often conflate the epistemic game $(G, \mathcal{T})$ with the associated type structure $\mathcal{T}$.

Note, an epistemic game $(G, \mathcal{T})$ induces a set of **states** $\prod_{i \in I}(S_i \times T_i)$. That is, a state describes a strategy-type pair for each player. We will focus on the state of states at which players are "rational and reason about rationality." This will be formalized as rationality and m^{th}-order belief of rationality.

Rationality is the concept that each player maximize their subjective expected payoffs, given their belief about how the game if played. Because a type specifies beliefs, this is a property of a strategy-type pair.

Definition 3.3. 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 3.4. 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.[6] Given some $E_{-i} \subseteq S_{-i} \times T_{-i}$, write

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

[5] Appendix E.1 extends the analysis to type structures that only induce finite-order beliefs.

[6] We use the phrase "event" for a measurable set.

for the set of types that believe E_{-i}. Note, if $E_{-i} = \emptyset$, then $B_i(E_{-i}) = \emptyset$. If E_{-i} is Borel then $B_i(E_{-i})$ is Borel. (See Lemma A.1.)

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 3.5. 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 $R^\infty = \prod_{i \in I} R_i^\infty$.

The following example illustrates these definitions.

Example 3.2. Refer back to the Undercutting Game with $X = 3$ and the type structure in Example 3.1. Given the first-order beliefs associated with $t_i \in \{t_i^1, t_i^2\}$, 1 is the unique best response; given the first-order beliefs associated with t_i^3, 2 is the unique best response. Thus,

$$R_i^1 = \{(1, t_i^1), (1, t_i^2), (2, t_i^3)\}.$$

Note, t_i^1 and t_i^2 both believe R_{-i}^1, but t_i^3 does not. Thus,

$$R_i^2 = \{(1, t_i^1), (1, t_i^2)\}.$$

With this, t_i^2 does not believe R_{-i}^2, but t_i^1 does. In fact, for each $m \geq 3$, $R_i^3 = \{(1, t_i^1)\}$. □

4 The Standard Benchmark Result

We begin with a baseline result.

Proposition 4.1. *Fix an epistemic game $(G, \mathcal{T})$ where $\mathcal{T}$ is complete. Then, for each m, $\mathrm{proj}_S R^m = S^m$.*

Quite generally, the behavior consistent with R$(m-1)$BR is m-rationalizable (i.e., $\mathrm{proj}_S R^m \subseteq S^m$). Proposition 4.1 says something more specific: In a complete type structure, the predictions of R$(m-1)$BR coincide with the m-rationalizable strategies (i.e., $\mathrm{proj}_S R^m = S^m$). (See, e.g., Proposition 1 in Friedenberg and Keisler, forthcoming.)

To better understand this point, it will be useful to return to the Undercutting Game. First observe that, for any type structure, $\mathrm{proj}_S R^m \subseteq S^m$. (This is true even if the type structure is not complete.) In particular, if (s_i, t_i) is rational, then s_i must be a best response under the first-order belief $\mathrm{marg}_{S_i} \beta_i(t_i)$. As such, s_i cannot be the dominated strategy X. As a consequence, if t_i believes that the other player is rational, then t_i must assign probability 1 to the other player choosing some strategy (strictly) less than X. Thus, if (s_i, t_i) is "rational and believes rationality," then s_i cannot be $X - 1$. And so on.

For the converse, fix a 1-rationalizable strategy s_i for i, i.e., a strategy (strictly) less than X. Note, s_i is a best response, if i believes that he is undercutting the other player by 1, i.e., if i believes the other player chooses the strategy $s_{-i} = s_i + 1$. If the type structure is complete, there exists some type t_i that holds that belief. For that type t_i, (s_i, t_i) is rational. Moreover, analogous arguments holds for higher levels of rationalizability.

Note, for the converse, it is important that the type structure is complete. If it is not, there may be m-rationalizable strategies that are inconsistent with R$(m-1)$BR. To see this, fix a type structure where each type t_i has a first-order belief that assigns probability 0 to $s_{-i} = X$. In that case, if (s_i, t_i) is rational then $s_i \leq X - 2$. That is, rationality precludes the 1-rationalizable strategy $X - 1$.

As a corollary to Equation (1) and Proposition 4.1, we have the following:

Corollary 4.1. *Fix an epistemic game $(G, \mathcal{T})$ where $\mathcal{T}$ is complete, and some $\mu = (\mu_i : i \in I) \in \prod_{i \in I} \Delta(S_{-i})$. Then, for each m,*

$$\bigcup_{k \geq m} L^k[\mu] \subseteq S^m = \mathrm{proj}_S R^m.$$

Note, since $\bigcup_{k \geq m} L^k[\mu]$ may be strictly contained in the m-rationalizable strategy set, it can also be strictly contained in the R$(m-1)$BR prediction.

5 Level-k Type Structures

The Undercutting Game illustrates that the level-k solution concept may preclude (at all levels) strategies that are consistent with rationality. To understand why, observe that a complete type structure allows players to hold any beliefs. (In fact, as we have seen, that is crucial from the perspective of Proposition 4.1.) On the other hand, the level-k analysis implicitly restricts the players' higher-order beliefs, so that they are generated by the first-order beliefs μ: An L1 player has a first-order belief μ_i. An L2 player assigns probability 1 to "other players have the first-order beliefs μ_{-i}." And so on.

This suggests that foundations for the level-k solution concept will require exogenously restricting the players' beliefs. We capture those exogenous restrictions by introducing the concept of a level-k type structure. To do so, it will be useful to have a definition: Given a type structure $\mathcal{T} = (S_{-i}, T_i, \beta_i : i \in I)$, call $\{T_i^k : k = 1, 2, \ldots\}$ a **Borel cover** of T_i if: (i) each T_i^k is a Borel subset of T_i, and (ii) $\bigcup_{k \geq 1} T_i^k = T_i$.[7]

Definition 5.1. Fix a $\mu = (\mu_i : i \in I) \in \prod_{i \in I} \Delta(S_{-i} \times T_{-i})$. Call an S-based type structure $\mathcal{T} = (S_{-i}, T_i, \beta_i : i \in I)$ a **level-k type structure for** μ if, for each i, there exists a Borel cover $\{T_i^k : k = 1, 2, \ldots\}$ of T_i so that the following hold:

(i) for each k, $T_i^k \neq \emptyset$,

(ii) for each $t_i \in T_i^1$, $\mathrm{marg}_{S_{-i}} \beta_i(t_i) = \mu_i$, and

(iii) for each k and each $t_i \in T_i^{k+1}$, $\beta_i(t_i)(S_{-i} \times T_{-i}^k) = 1$.

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

Fix a level-k type structure for μ. For each k, this structure has a non-empty set of **k-types**, viz. T_i^k. The 1-types have the first-order beliefs associated with μ. The 2-types assign probability 1 to the 1-types having the first-order beliefs associated with μ. More generally, the $(k+1)$-types assign probability 1 to the k-types. Also, each type in the level-k type structure is a k-type for some k. (This corresponds to the requirement that $\{T_i^k : k = 1, 2, \ldots\}$ forms a Borel cover of T_i.) Thus, each type in the level-k type structure can be viewed as having a belief generated by the first-order beliefs μ.

The next series of examples point to some subtleties in level-k type structures and in k-types.

[7] An example of a Borel cover is a partition with measurable partition members. However, unlike a partition, a Borel cover does not require that any two members of the cover are disjoint.

Example 5.1. This example highlights the fact that the exogenously specified μ limits—but does not pin down—the hierarchies of beliefs in a level-k type structure. In particular, we will construct two S-based level-k type structures for μ, viz. $\mathcal{T} = (S_{-i}, T_i, \beta_i : i \in I)$ and $\mathcal{T}' = (S_{-i}, T'_i, \beta'_i : i \in I)$, that induce distinct hierarchies of beliefs. For each i, take $T_i = T'_i = \mathbb{N}_+$. Likewise, for each i, choose $\beta_i(1) = \beta'_i(1)$ so that $\text{marg}_{S_{-i}}\beta_i(1) = \mu_i$. Choose $\beta_i(2) \neq \beta'_i(2)$, but $\beta_i(2)(S_{-i} \times \{1\}) = \beta'_i(2)(S_{-i} \times \{1\}) = 1$. Finally, for each $k \geq 3$, choose $\beta_i(k) = \beta'_i(k)$ so that $\beta_i(k)(S_{-i} \times \{k-1\}) = 1$. It follows that both $\mathcal{T}$ and $\mathcal{T}'$ are level-k type structures for μ.

Note, $\text{marg}_{S_{-i}}\beta_i(2)(S_{-i} \times \{1\}) = \text{marg}_{S_{-i}}\beta'_i(2)(S_{-i} \times \{1\}) = 1$. So, the types $t_i = 2$ and $t'_i = 2$ have different first-order beliefs in the type structures $\mathcal{T}$ and $\mathcal{T}'$. Thus, in these type structures, type $t_i = 2$ and $t'_i = 2$ also have different second-order beliefs—i.e., beliefs about both the strategies and first-order beliefs of other players. That said, they do both assign probability 1 to the other players having first-order beliefs μ_{-i}.

Types $u_i = 3$ and $u'_i = 3$ have the same first-order beliefs in $\mathcal{T}$ and $\mathcal{T}'$. However, because $t_i = 2$ and $t'_i = 2$ have different first-order beliefs in these type structures, types $u_i = 3$ and $u'_i = 3$ have different second-order beliefs in these type structures. And so on. $\square$

Example 5.2. This example shows that the cover may not be unique. As a result, a type may be a k-type for one associated cover and an ℓ-type for another associated cover, despite the fact that $k \neq \ell$. Thus, the choice of terminology "k-type" is associated with a particular 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, for each $k \geq 1$, $\beta_i(k+1)(S_{-i} \times \{k\}) = 1$.

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

Example 5.3. This example shows that, for a given level-k type structure, we may not be able to choose the cover to be a partition. As such, we may have that a type is both a k-type and an ℓ-type for every associated cover.

Construct an S-based level-k type structures for μ, viz. $\mathcal{T} = (S_{-i}, T_i, \beta_i : i \in I)$, as follows: For each i, take $T_i = \mathbb{N}_+$. Choose β_i so that it satisfies the following properties. First, $\text{marg}_{S_{-i}}\beta_i(t_i) = \mu_i$ 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 $\{T_i^k : k = 1, 2, \ldots\}$ so that $T_i^1 = \{1, 3\}$ and, for each $k \geq 2$, $T_i^k = \{k\}$. This cover is non-partitional. However, any cover must be non-partitional. To see this, fix a cover $\{U_i^k : k = 1, 2, \ldots\}$. Since $\text{Supp}\,\text{marg}_{T_{-i}}\beta_i(1) = T_{-i}$, 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 $U_i^1 \cap U_i^3 \neq \emptyset$. $\square$

Examples 5.1-5.3 construct specific level-k type structures. The constructed type structures are sparse in the sense that they do not induce all possible hierarchies of beliefs consistent with the given first-order beliefs. We will also be interested in level-k type structures that are rich in the sense that they induce all possible hierarchies of beliefs that are consistent with the given first-order beliefs.

Definition 5.2. Fix a $\mu = (\mu_i : i \in I) \in \prod_{i \in I} \Delta(S_{-i} \times T_{-i})$. Call an S-based type structure $\mathcal{T} = (S_{-i}, T_i, \beta_i : i \in I)$ a **complete level-k type structure for** μ if, for each i, there exists a Borel cover $\{T_i^k : k = 1, 2, \ldots\}$ of T_i so that the following holds:

(i) for each k, $T_i^k \neq \emptyset$,

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

(iii) for each k and each $t_i \in T_i^{k+1}$, $\beta_i(t_i)(S_{-i} \times T_{-i}^k) = 1$, and

(iv) for each k and each $\nu_i \in \Delta(S_{-i} \times T_{-i})$ with $\nu_i(S_{-i} \times T_{-i}^k) = 1$, there exists $t_i \in T_i^{k+1}$ with $\beta_i(t_i) = \nu_i$.

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

So, $\mathcal{T}$ is a complete level-k type structure for μ if it is a level-k type structure for μ where we can find a cover that satisfies the following additional requirement: For each belief that assigns probability 1 to the k-types, there is a type of the player that holds that belief. In a sense, a complete level-k type structure is one in which the only exogenous restrictions on the players hierarchies of beliefs are those induced by the anchor. Incomplete level-k type structures can involve additional exogenous restrictions on the players' hierarchies of beliefs.

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

Proposition 5.1. *For each $\mu = (\mu_i : i \in I)$, there exists a complete level-k type structure for μ.*

It is useful to note that, for any complete level-k type structure, there is some associated Borel cover that is not a partition. To see this fix some complete level-k type structure for $\mu = (\mu_i : i \in I)$. Choose the Borel cover so that $T_i^1 = \{t_i \in T_i : \text{marg}_{S_{-i}} \beta_i(t_i) = \mu_i\}$ and, for each $k \geq 1$, $T_i^{k+1} = \{t_i \in T_i : \beta_i(t_i)(S_{-i} \times T_{-i}^k) = 1\}$. Now observe that, for each $k \geq 2$, there must be some t_i^k with $\text{marg}_{S_{-i}} \beta_i(t_i^k) = \mu_i$. (In fact, there will be many such types.) Thus, for each $k \geq 2$, $T_i^k \cap T_i^1 \neq \emptyset$. With this, for each $k \geq 3$, there exists some type u_i^k so that $\beta_i(u_i^k)(T_{-i}^{k-1} \cap T_{-i}^1) = 1$. Any such $u_i^k \in T_i^k \cap T_i^2$. Thus, for each $k \geq 3$, $T_i^k \cap T_i^2 \neq \emptyset$. And so on, inductively.

The proof of Proposition 5.1 constructs a particular complete level-k type structure for μ—one that is, in a sense, canonical. For that type structure, the only associated Borel cover is non-partitional. This occurs for a similar reason that the cover in Example 5.3 cannot be a partition. See Remark B.2.

6 Toward Foundations for Level-k Behavior

Recall, the k-types have beliefs that are, in a sense, determined by the exogenous first-order belief μ. As such, there is a natural intuition that, for k-types, the $R(k-1)$BR predictions should coincide with level-k behavior $L^k[\mu]$. However, this intuition is not quite right. To understand why, we begin with a preliminary result.

Proposition 6.1. *Fix an epistemic game $(G, \mathcal{T})$ where $\mathcal{T}$ is a level-k type structure for μ. Then:*

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

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

Proposition 6.1 has two parts. First, it says that the set of rational predictions for 1-types coincide with the level-1 behavior. This corresponds to the standard intuition. Second, it says that, for each k-type, any prediction of R$(k-1)$BR is *consistent* with level-k behavior. It stops short of saying that such predictions *coincide* with level-k behavior. In fact, they may not. We may well have that $\text{proj}_{S_i}\left(R_i^k \cap (S_i \times T_i^k)\right) \subsetneq L_i^k[\mu]$. The next two examples illustrate two reasons that the inclusion may be strict.

Example 6.1. Return to the Undercutting Game in Figure 2.1, but make one change: Now, take $\mu_1(X) = \mu_2(X) = 1$. So, now, the anchor assigns probability 1 to the other player choosing the highest action.

Construct a level-k type structure 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 $\beta_i(1)$ so that $\text{Supp}\,\beta_i(1) = \{X\} \times T_{-i}$. For each $k \geq 2$, choose $\beta_i(k)$ to satisfy the following criteria. First, $\beta_i(k)(S_{-i} \times \{k-1\}) = 1$. Second, $\text{marg}_{S_2}\beta_1(2)(X) = 1$ but $\text{marg}_{S_1}\beta_2(2)(X) \neq 1$. Third, there is no i and $k \geq 3$ with $\text{marg}_{S_{-i}}\beta_i(k) = \mu_i$.

This is a level-k type structure for μ. But the only cover has $T_1^2 = \{2\}$. Moreover, $t_1 = 2$ does not believe R_2^1, since X is dominated. Thus, $R_1^2 \cap (S_1 \times T_1^2) = \emptyset$. □

Example 6.2. Figure 6.1 depicts a revised undercutting game. Now, if the other player chooses a number $s_{-i} \geq 2$, player i has an incentive to undercut the other player by either 1 or 2. (As previously, if the other player chooses the lowest number $s_{-i} = 1$, then i has an incentive to match the other player.)

	1	2	3	$\cdots$	$X-1$	X
1	1,1	1,0	1,0	$\cdots$	0,0	0,0
2	0,1	0,0	1,0	$\cdots$	0,0	0,0
3	0,1	0,1	0,0	$\cdots$	0,0	0,0
4	0,0	0,1	0,1	$\cdots$	0,0	0,0
$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$
$X-3$	0,0	0,0	0,0	$\cdots$	1,0	0,0
$X-2$	0,0	0,0	0,0	$\cdots$	1,0	1,0
$X-1$	0,0	0,0	0,0	$\cdots$	0,0	1,0
X	0,0	0,0	0,0	$\cdots$	0,1	0,0

Figure 6.1: Revised Undercutting Game

Set $\mu_1(X) = \mu_2(X) = 1$. So, the anchor assigns probability 1 to the other player choosing the highest action. As such, $L_i^1[\mu] = \{X-1, X-2\}$ and $L_i^2[\mu] = \{X-2, X-3\}$. More generally,

$$L_i^k[\mu] = \begin{cases} \{X-k, X-k-1\} & \text{if } X-2 \geq k \\ \{X-k\} & \text{if } k = X-1 \\ \{1\} & \text{if } k \geq X. \end{cases}$$

Construct an level-k type structure 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 each β_i so that, for each $k \geq 1$: (i) $\text{marg}_{S_{-i}}\beta_i(k)(X-k+1) = 1$, and (ii) $\beta_i(k+1)(S_{-i} \times \{k\}) = 1$. Observe that $\text{marg}_{S_{-i}}\beta_i(1) = \mu_i$.

This is a level-k type structure for μ. But, there is only one cover and, for that cover, each $T_i^k = \{k\}$. As such, $(s_i, t_i) \in R_i^2 \cap (S_i \times T_i^2)$ if and only if $s_i = X-2$. However, $L_i^2[\mu] = \{X-2, X-3\}$. Thus, $\text{proj}_{S_i}(R_i^2 \cap (S_i \times T_i^2)) \subsetneq L_i^2$. □

The examples illustrate two reasons why the inclusion may be strict. First, $R_i^k \cap (S_i \times T_i^k)$ may be empty. This can arise because a k-type's first-order beliefs may be such that it cannot believe $R_{-i}^{k-1}[\mu]$. Second, even if $R_i^k \cap (S_i \times T_i^k)$ is non-empty, it may still be strictly contained in $L_i^k[\mu]$. This can only arise if $L_i^{k-1}[\mu]$ is not a singleton.[8] In that case, T_i^k may simply not induce all the first-order beliefs that are consistent with L_i^{k-1}.

Both of these phenomena can arise because an arbitrary level-k type structure permits a sparse set of beliefs. Neither phenomena can arise in a complete level-k type structure—since, there, for each $k \geq 2$, the sets T_i^k are "rich."

Theorem 6.1. *Fix an epistemic game $(G, \mathcal{T})$, where $\mathcal{T}$ is a complete level-k type structure for μ. Then, for each k,*

$$\text{proj}_{S_i} \left(R_i^k \cap (S_i \times T_i^k)\right) = L_i^k[\mu].$$

Theorem 6.1 says that, in any complete level-k type structure, the R$(k-1)$BR predictions associated with k-types coincides (exactly) with the level-k behavior. This captures a standard rationale behind the level-k model: If a player's beliefs about beliefs are determined by k steps of reasoning about an anchor μ and, if they satisfy R$(k-1)$BR, then their behavior is level-k behavior for μ. And, conversely, level-k behavior is consistent with reasoning in this way.

7 Reasoning about Rationality in Level-k Type Structures

Theorem 6.1 captures the standard intuition associated with level-k models. One might be tempted to conclude that—by restricting attention to level-k type structures—there is an equivalence between reasoning about rationality (in the sense of RmBR) and the level-k solution concept. That is, drawing from Proposition 4.1, one might conjecture that, in a complete level-k type structure,

$$\bigcup_{k \geq m} L^k[\mu] = \text{proj}_S R^m \subseteq S^m.$$

However, this is incorrect. The key is the following:

Theorem 7.1. *Fix an epistemic game $(G, \mathcal{T})$, where $\mathcal{T}$ is a complete level-k type structure for μ. Then, for each m, $\text{proj}_S R^m = S^m$.*

To understand Theorem 7.1, fix a complete level-k type structure. Theorem 6.1 tells us that: (i) for each strategy $s_i \in L_i^1[\mu]$ there is a 1-type t_i^1 so that (s_i, t_i^1) is rational, and (ii) for each 1-type t_i^1 and (s_i, t_i^1) rational, $s_i \in L_i^1[\mu]$. So it establishes an equivalence between the level-1 behavior and the behavior of 1-types that are rational. However, it is silent about the behavior of 2-types that are rational. In fact, because the type structure is a complete level-k type structure, for each first-order belief that a player can hold, there is a 2-type of the player that holds that first-order belief. As such, for any 1-rationalizable strategy s_i there is a 2-type t_i^2 so that (s_i, t_i^2) is rational. Thus, if we look at all the rational strategy-type pairs—not just those associated with 1-types—our behavioral prediction is the set of 1-rationalizable strategies. The following example illustrates this point.

[8] The non-singleton case is *not* simply a thought exercise. In particular, it arises naturally in the N-player Beauty Contest game (Nagel, 1995), provided $N \geq 3$.

Example 7.1. Return to the Undercutting Game. As in Section 2, take $\mu_1(\mathbb{E}(x)) = \mu_2(\mathbb{E}(x)) = 1$. Fix a complete level-$k$ type structure for $\mu = (\mu_1, \mu_2)$. We will argue that, in that type structure, $\text{proj}_{S_i} R_i^m = S_i^m$ for each $m \geq 1$.

Recall that, in each type structure, $\text{proj}_{S_i} R_i^m \subseteq S_i^m$. So we focus on showing that $S_i^m \subseteq \text{proj}_{S_i} R_i^m$. To do so, first, fix some $s_i \in S_i^1 = \{1, 2, \ldots, X-1\}$. Note, there is some 2-type $t_i^2[s_i] \in T_i^2$ so that $\beta_i(t_i^2)(\{(s_i+1)\} \times T_{-i}^1) = 1$. For such a type $t_i^2[s_i]$, we have $(s_i, t_i^2[s_i]) \in R_i^1$. Next, fix some $s_i \in S_i^2 = \{1, 2, \ldots, X-2\}$. Then there exists some 3-type $t_i^3[s_i] \in T_i^3$ so that $\beta_i(t_i^3[s_i])(\{(s_i+1)\} \times t_{-i}^2[(s_i+1)]) = 1$. Since $((s_i+1), t_{-i}^2[s_i+1]) \in R_i^2$, it follows that $(s_i, t_i^3[s_i]) \in R_i^2$.

We can continue inductively. In particular, for each $m \leq X-1$, the following holds: If s_i is an m-rationalizable strategy s_i, we can find an $(m+1)$-type $t_i^{m+1} \in T_i^{m+1}$ so that (s_i, t_i^{m+1}) is consistent with R$(m-1)$BR. The type t_i^{m+1} will assign probability 1 to the other player playing $s_{-i} = (s_i+1) \in S_{-i}^m$ and the other player having an m-type t_{-i}^m with $(s_{-i}, t_{-i}^m) \in R_{-i}^{m-1}$. $\square$

As a consequence of Theorem 7.1, we have the following corollary:

Corollary 7.1. *Fix an epistemic game $(G, \mathcal{T})$, where $\mathcal{T}$ that is a complete level-k type structure for μ.*

(i) For each $m \geq 1$, $\bigcup_{k \geq m} L^k[\mu] \subseteq S^m = \text{proj}_S R^m$.

(ii) If each $\mu_i(S_{-i}^\infty) = 1$, then, for each $m \geq 1$, $\bigcup_{k \geq 1} L^k[\mu] \subseteq S^m = \text{proj}_S R^m$.

Corollary 7.1 says that, for each $k \geq m$, the level-k behavior is consistent with R$(m-1)$BR in a complete level-k type structure. Moreover, if the anchor assigns probability 1 to the rationalizable strategies, then, for each k and each m, the level-k behavior is consistent with R$(m-1)$BR. Note, even in this last case, the inclusion may be strict.[9]

8 Discussion

A. Foundations for the Level-k Solution Concept. Theorem 7.1 states that, in a complete level-k type structure, the RmBR behavior corresponds to the $(m+1)$-rationalizable strategies. By contrast, Theorem 6.1 states that, in a complete level-k type structure, the RmBR behavior for $(m+1)$-types corresponds to level-$(m+1)$ behavior. Thus, if we associate k-types with the epistemic conditions of R$(k-1)$BR, we can justify level-k behavior.

Note, this justification of level-k behavior rests on assigning different hierarchies of beliefs with different epistemic conditions—a step that is atypical in the epistemic game theory literature.[10] Importantly, for certain types, there will be no way to uniquely assign types to a single epistemic condition. This is because it may not be possible to associate a complete level-k type structure with a partition. (Refer to the discussion immediately following Proposition 5.1.) And even if it were possible to associate a complete level-k type structure with some partition, there may be another cover that is non-partitional. So, arguably, the result cannot be seen as endogenizing the epistemic conditions.

In light of this fact, we conclude that the approach to iterated reasoning taken in the level-k literature is quite different from the approach taken in the epistemic game theory literature.

[9] A prominent example of this is Alaoui and Penta's (2014) version of the 11-20 game: Take $\mu_1(20) = \mu_2(20) = 1$ and observe that the anchor assigns probability 1 to a rationalizable strategy. Then the union of the level-k behavior is $\{11, 12, \ldots, 19\}$, but the entire strategy set is rationalizable.

[10] There is good reason for that: By allowing oneself to tailor the epistemic conditions to the hierarchies, the researcher may well end up justifying any solution concept. Restricting the ability to tailor the epistemic conditions to the players beliefs imposes discipline on the epistemic analysis.

B. Hierarchies of Beliefs vs. Finite-Order Beliefs We made use of an epistemic type structure that induces hierarchies of beliefs about the strategies played. Such hierarchies of beliefs require unlimited ability to engage in interactive reasoning—i.e., the ability to specify all sentences of the form "I think that you think that I think ...." However, often, the level-k solution concept is motivated by a stipulation that players have a limited ability to engage in such sentences.

We can recast the analysis here in terms of an epistemic model where epistemic types only induce finite-order beliefs (as in Heifetz and Kets, 2013 or Kets, 2010). The key is that the epistemic conditions of RmBR depend only on the $(m+1)^{\text{th}}$-order beliefs. As such, from the perspective of Theorems 6.1-7.1, we can either use the standard definition of a type structure—one that induces full hierarchies of beliefs—or we can use a modified definition of a type structure—one that induces finite-order beliefs but (potentially) not the full hierarchy of beliefs.

Appendix E makes this idea precise. We begin with an extended definition of an epistemic type structure, where takes can now induce finite-order beliefs but may not induce the full hierarchy. (An ordinary epistemic type structure is a special case of this type structure.) We focus on the special case where the game has no weakly dominant strategy and the extended type structure induces all first-order beliefs. Proposition E.1 shows that we can convert this type structure into an ordinary type structure in a way that preserves RmBR behavior. Thus, in this case, we can make use of an ordinary type structure—instead of such an extended type structure—from the perspective of analyzing RmBR (as it is used in this paper).

The requirement that the extended type structure induces all first-order beliefs may, at first, appear restrictive. However, note, it is a condition that is satisfied by both an ordinary complete and an ordinary complete level-k type structure. So, arguably, it is an assumption that is desirable for the purpose of the exercise in this paper.

C. Alternate "Rich" Level-k Type Structures: Taken together, Examples 6.1-6.2 and Theorem 7.1 point to a tension. If we allow the level-k type structure to be sparse, then R$(m-1)$BR may preclude strategies allowed by the level-m behavior. On the other hand, if we require the level-k type structure to be rich (in the sense of completeness), then R$(m-1)$BR is characterized by m-rationalizability and *not* by the level-k solution concept. This raises the question of whether there are intermediate notions of "richness" that would allow R$(m-1)$BR to be characterized by level-k behavior for $k \geq m$.

First, note that we can always rig a type structure to deliver the desired output. In particular, we can construct the first-order beliefs associated with each k-type, by making reference to the level-k solution concept. As such, the very notion of the type structure would make reference to the solution concept for which we are attempting to provide foundations. From the perspective of epistemic game theory, this is undesirable. If we allow the epistemic assumptions to depend on details of the solution concept, then the epistemic analysis cannot illuminate the extent to which "reasoning" gives rise to the solution concept. See Chapter 6-Section 6.5 in Battigalli, Friedenberg and Siniscalchi, 2012.

But, there is a potential second approach—e.g., requiring that the notion of "richness" only refers to marginal beliefs. This would be a modification of Condition (iv) in Definition 5.2. There are two obvious possibilities:

(iv') for each k and each $\nu_i \in \Delta(T_{-I})$ with $\nu_I(T^k_{-i}) = 1$, there exists $t_i \in T_i^{k+1}$ with $\operatorname{marg}_{T_{-i}} \beta_i(t_i) = \nu_i$;

or

(iv”) for each k and each $\nu_i \in \Delta(S_{-i})$, there exists $t_i \in T_i^{k+1}$ with $\text{marg}_{S_{-i}}\beta_i(t_i) = \nu_i$.

However, neither definition suffices. If we replace condition (iv) by either (iv’) or (iv”), we may still have that $\text{proj}_{S_i}(R_i^k \cap (S_i \times T_i^k)) \subsetneq L_i^k$. Example 6.2 shows this for (iv’). Appendix E shows this for (iv”).

D. RmBR Behavior of k-Types Let us revisit how the R$(m-1)$BR behavior of k-types relates to the R$(m-1)$BR behavior. In a level-k type structure for μ, we have

$$\bigcup_{k \geq m} \text{proj}_{S_i} \left(R_i^m \cap (S_i \times T_i^k)\right) \subseteq \text{proj}_{S_i} R_i^m.$$

Thus, for each $k \geq m$, the R$(m-1)$BR behavior of k-types is contained in the R$(m-1)$BR behavior. However, when each μ_i assigns positive probability to a dominated strategy, we can draw a tighter conclusion.

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

$$\bigcup_{k \geq m} \text{proj}_{S_i} \left(R_i^m \cap (S_i \times T_i^k)\right) = \text{proj}_{S_i} R_i^m.$$

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

E. Identification of Reasoning about Rationality The results here speak to identification of reasoning about rationality—in the sense of RmBR—for the case where the players’ type structure is entirely determined by the anchor. Importantly, the conclusions are different from the identification of reasoning often inferred in the level-k literature. To see this, consider two examples: the 11-20 game (Arad and Rubinstein, 2012) and the Beauty Contest game (Nagel, 1995).

First, consider the 11-20 game. Two players simultaneously choose a number in $\{11, 12, \ldots, 20\}$. If Player i chooses $s_i = s_{-i} - 1$, her payoff is $s_i + 20$; if Player i chooses any other s_i, her payoff is s_i. Note the entire strategy set is rationalizable. (Each $s_i \leq 19$ is a best response to $s_i + 1$; moreover, $s_i = 20$ is a best response to $s_{-i} = 11$.) Thus, any anchor assigns probability 1 to a rationalizable strategy. As such, for any anchor—a fortiori the (typical) anchor that assigns probability one to 20—and any k, the level-k behavior is consistent with RmBR for all m—even if we restrict attention to type structures defined (only) by the anchor.

Second, consider a parameterized Beauty Contest game. Three players simultaneously choose a number in $\{1, 2, 3, 4, 5\}$. If players choose (s_1, s_2, s_3), then the player closest to $(^2/_3)$s of the average $(^{s_1+s_2+s_3}/_3)$ has a payoff of 1 and the other players have a payoff of 0. (If players tie others as being ‘closest to the average,’ they split the payoff of 1 equally.) Note, 5 is dominated by any mixture on $\{1, 2\}$; all other strategies s_i are a best response to $s_{-i} = (s_i + 1, s_i + 1)$. In fact, $S_i^1 = \{1, 2, 3, 4\}$, $S_i^2 = \{1, 2, 3\}$, $S_i^3 = \{1, 2\}$, and, for each $m \geq 4$, $S_i^m = \{1\}$.

Suppose the anchor assigns equal weight to all strategy profiles s_{-i}. Note, in this case, $\mu_i(S_{-i}\backslash S^1_{-i}) > 0$. It is easily verified that $L^1_i[\mu] = \{2\}$ and, for each $k \geq 2$, $L^k_i[\mu] = \{1\}$. Thus, if we observe either of $3, 4$, it would be interpreted as play of a level-0 player, even though 4 is consistent with rationality (resp. 3 is consistent with R1BR) when we restrict attention to type structures defined (only) by the anchor. Likewise, if we observe 2, it would be interpreted as play of a level-1 player, even though the behavior is consistent with R2BR when we restrict attention to type structures defined (only) by the anchor.

Appendix A Proofs for Sections 5

Lemma A.1.

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

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

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 A.2. *Let Ω_1, Ω_2 be metrizable spaces where $|\Omega_1| < \infty$. Then the mapping* $\text{marg}_{\Omega_1} : \Delta(\Omega_1 \times \Omega_2) \to \Delta(\Omega_1)$ *is continuous.*

Proof. Fix a sequence $(\nu^n : n = 1, 2, \ldots)$ where each $\nu^n \in \Delta(\Omega_1 \times \Omega_2)$. Suppose $(\nu^n : n = 1, 2, \ldots)$ converges to ν^*. Write $\hat{\nu}^n = \text{marg}_{\Omega_1}(\nu^n)$ and $\hat{\nu}^\infty = \text{marg}_{\Omega_1}(\nu^\infty)$. It suffices to show that $(\hat{\nu}^n : n = 1, 2, \ldots)$ converge to $\hat{\nu}^\infty$. For that, it suffices to show that, for each $\omega_1 \in \Omega_1$, $\lim_{n\to\infty} \hat{\nu}^n(\{\omega_1\}) = \hat{\nu}^\infty(\{\omega_1\})$ or, equivalently, $\lim_{n\to\infty} \nu^n(\{\omega_1\} \times \Omega_2) = \nu^\infty(\{\omega_1\} \times \Omega_2)$.

Fix some $\{\omega_1\} \times \Omega_2$ and observe that this set is clopen. As such, $\partial(\{\omega_1\} \times \Omega_2) = \emptyset$. So, by Theorem 15.3 in Aliprantis and Border (2007), $\lim_{n\to\infty} \nu^n(\{\omega_1\} \times \Omega_2) = \nu^\infty(\{\omega_1\} \times \Omega_2)$. ■

Lemma A.3. *For each m, the sets R^m_i are Borel.*

Proof. The proof is by induction on m.

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

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

Using Berge's Theorem and the fact that $\Delta(S_{-i})$ is metrizable, this set is closed. So, by Lemma A.2,

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

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

$$R^1_i = \bigcup_{s_i \in S_i} \left(\{s_i\} \times \beta_i^{-1}(\{\hat{P}[s_i]\})\right)$$

and, therefore, R^1_i is Borel.

$m \geq 2$: Assume that, for each i, R_i^m is Borel. As such, each R_{-i}^m is also Borel. So by Lemma A.1(i), R_i^m is Borel. ∎

Appendix B Proof of Proposition 5.1

This appendix proves Proposition 5.1. Before coming to the proof, it will be useful to introduce some mathematical preliminaries: Given a Borel subset of Ω, viz. $\Phi \subseteq \Omega$, endow Φ with the relative topology. If $\nu \in \Delta(\Omega)$ with $\nu(\Phi) = 1$, we can construct a $\tilde{\nu} \in \Delta(\Phi)$ so that, for each Borel $E \subseteq \Phi$, $\tilde{\nu}(E) = \nu(E)$. Note, $\tilde{\nu}$ is indeed a probability measure since $\nu(\Phi) = 1$. Call $\tilde{\nu}$ the **restriction of ν to Φ**.

We now turn to the proof. Fix some $\mu = (\mu_i : i \in I) \in \prod_{i \in I} \Delta(S_{-i})$. Fix, also, a complete S-based type structure $\mathcal{T}^* = (S_{-i}, T_i^*, \beta_i^* : i \in I)$. We will inductively define sets $E_i^k \subseteq T_i^*$. Set

$$E_i^1 = \{t_i^* \in T_i^* : \operatorname{marg}_{S_{-i}} \beta_i^*(t_i^*) = \mu_i\}.$$

Assuming E_{-i}^k has been defined, set $E_i^{k+1} = B_i(S_{-i} \times E_{-i}^k)$. Then set $E_i^\infty = \bigcup_{k \geq 1} E_i^k$. It will also be convenient to define $\hat{E}_i^1$ to be the set of $t_i^* \in E_i^1$ so that $\beta_i^*(t_i^*)(S_{-i} \times E_{-i}^\infty) = 1$.

Lemma B.1.

(i) For each k, E_i^k is Borel and non-empty.

(ii) The set E_i^∞ is Borel and non-empty.

(iii) The set $\hat{E}_i^1$ is Borel and non-empty.

Proof. We begin by proving part (i). Part (i) is proved by induction on k.

Begin with $k = 1$. Note

$$\{\nu_i \in \Delta(S_{-i} \times T_{-i}^*) : \nu_i = \mu_i\}$$

is nonempty and Borel. (See Lemma A.2.) Since β_i^* is measurable, E_i^1 is Borel. By completeness, $E_i^1 \neq \emptyset$.

Assume that the claim holds for k, then the set

$$\{\nu_i \in \Delta(S_{-i} \times T_{-i}^*) : \nu_i(S_{-i} \times E_{-i}^k) = 1\}$$

is non-empty and, by Aliprantis and Border (2007, Lemma 15.16), it is Borel. Since β_i^* is measurable, E_i^{k+1} is Borel. By completeness, $E_i^{k+1} \neq \emptyset$. This establishes part (i).

Part (ii) follows immediately from part (i). From that

$$\{\nu_i \in \Delta(S_{-i} \times T_{-i}^*) : \nu_i(S_{-i} \times E_{-i}^\infty) = 1\}$$

is non-empty and Borel (Aliprantis and Border, 2007, Lemma 15.16). Since β_i^* is measurable, $\hat{E}_i^1$ is Borel. And, by completeness, $\hat{E}_i^1 \neq \emptyset$, which establishes part (iii). ∎

Let $T_i^{**} = E_i^\infty$ and endow T_i^{**} with the relative topology. Note, the Borel sets in T_i^{**} are also Borel sets in T_i^*. (See Aliprantis and Border, 2007, Lemma 4.20.) Now define β_i as follows: For each $t_i \in \hat{E}_i^1 \cup \bigcup_{k \geq 2} E_i^k$, let $\beta_i^{**}(t_i) \in \Delta(S_{-i} \times T_{-i}^{**})$ be the restriction of $\beta_i^*(t_i)$ to $S_{-i} \times T_{-i}^{**}$. (This is well defined since each such $\beta_i^*(t_i)$ assigns probability 1 to $S_{-i} \times T_{-i}^{**}$.) Now choose some $u_{-i}^* \in T_{-i}^{**}$. For each $t_i \in E_i^1 \backslash \hat{E}_i^1$, let $\beta_i^{**}(t_i) \in \Delta(S_{-i} \times T_{-i}^{**})$ be the unique measure $\mu_i^* \in \Delta(S_{-i} \times T_{-i}^{**})$ so that (i) $\text{marg}_{S_{-i}} \mu_i^* = \mu_i$, and (ii) $\mu_i^*(S_{-i} \times \{u_{-i}^*\}) = 1$.

Let $\text{id}_{-i} : S_{-i} \times T_{-i}^{**} \to S_{-i} \times T_{-i}^*$ be the identity map and note that id_{-i} is bimeasurable. Let $\underline{\text{id}}_{-i} : \Delta(S_{-i} \times T_{-i}^{**}) \to \Delta(S_{-i} \times T_{-i}^*)$ map each $\nu \in \Delta(S_{-i} \times T_{-i}^{**})$ into the image measure of ν under id_{-i}. Observe that, for each such $t_i \in \hat{E}_i^1 \cup \bigcup_{k \geq 2} E_i^k$, $\beta_i^*(t_i)$ is the image measure of $\beta_i^{**}(t_i)$ under id_{-i}.

Lemma B.2. *The mapping β_i^{**} is measurable.*

Proof. Fix some F_{-i} that is Borel in $\Delta(S_{-i} \times T_{-i}^{**})$. First suppose that $\mu_i^* \notin F_{-i}$. By construction, $(\beta_i^{**})^{-1}(F_{-i}) = T_i^{**} \cap (\beta_i^*)^{-1}(\underline{\text{id}}_{-i}(F_{-i}))$. Note that $\underline{\text{id}}_{-i}(F_{-i})$ is Borel. (See Aliprantis and Border, 2007, Lemma 15.4.) Thus, $(\beta_i^{**})^{-1}(F_{-i})$ is the intersection of two Borel sets and therefore Borel.

Next suppose that $\mu_i^* \notin F_{-i}$. By construction, $(\beta_i^{**})^{-1}(F_{-i}) = (T_i^{**} \cap (\beta_i^*)^{-1}(\underline{\text{id}}_{-i}(F_{-i}))) \cup (E_i^1 \backslash \hat{E}_i^1))$. Repeating the above argument and noting that $(E_i^1 \backslash \hat{E}_i^1)$ is Borel, we conclude $(\beta_i^{**})^{-1}(F_{-i})$ is Borel. ■

Proof of Proposition 5.1. We will show that the S-based type structure $\mathcal{T}^{**} = (S_{-i}, T_i^{**}, \beta_i^{**} : i \in I)$ is a complete level-k type structure for μ.

For each k, let $T_i^k = E_i^k$. Note that $\{T_i^k : k = 1, 2, \ldots\}$ forms a Borel cover of T_i^{**} with each $T_i^k \neq \emptyset$. Moreover, by construction, for each $t_i \in T_i^1 = E_i^1$,

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

Likewise, for each $t_i \in T_i^{k+1} = E_i^{k+1}$, $\beta_i^*(t_i)(S_{-i} \times E_{-i}^k) = 1$. As such, $\beta_i^{**}(t_i)(S_{-i} \times T_{-i}^k) = 1$.

Finally, fix some $\nu_i \in \Delta(S_{-i} \times T_{-i}^{**})$ with $\nu_i(S_{-i} \times T_{-i}^k) = 1$. Then there exists some $\nu_i^* \in \Delta(S_{-i} \times T_{-i}^*)$ so that ν_i is the image measure of ν_i^* under id_{-i}. By completeness, there exists a type $t_i^* \in T_i^*$ so that $\beta_i^*(t_i^*) = \nu_i^*$. Since $E_{-i}^k = T_{-i}^k$, by construction $t_i^* \in E_i^{k+1} = T_i^{k+1}$ and $\beta_i^{**}(t_i^*) = \nu_i$. ■

Remark B.1. The proof of Proposition 5.1 constructs a level-k type structure for μ associated with associated covers $\{T_i^k : k = 1, 2, \ldots\}$ that are non-partitional. In fact, for each $k \geq 1$, $T_i^1 \cap T_i^{k+1} \neq \emptyset$. To see this, recall that $\mathcal{T}^*$ is complete. Thus, for each $k \geq 1$, there is some type $t_i^* \in T_i^*$ with both $\text{marg}_{S_{-i}} \beta_i^{**}(t_i^*) = \mu_i$ and $\beta_i^{**}(t_i^*)(S_{-i} \times T_{-i}^k) = 1$. Under the associated cover, $t_i^* \in T_i^1 \cap T_i^{k+1}$. □

Remark B.2. We can push the point in Remark B.1 further: If $\left(\{U_i^k : k = 1, 2, \ldots\} : i \in I\right)$ is a collection of covers that jointly satisfy conditions (i)–(iv) for the type structure constructed in Proposition 5.1 then some $\{U_i^k : k = 1, 2, \ldots\}$ is non-partitional.

To see this, suppose otherwise, i.e., we have a collection of covers $\left(\{U_i^k : k = 1, 2, \ldots\} : i \in I\right)$ that are partitional and jointly satisfy conditions (i)–(iv) for the type structure constructed in Proposition 5.1. There are two cases to consider.

Case 1: $T_i^k = U_i^k$ for each i and each $k \geq 2$. In this case, each $U_i^1 \subsetneq T_i^1$. Note, there is some $t_i \in T_i^2$ with $\beta_i^{**}(t_i)(S_{-i} \times T_i^1 \backslash U_i^1) = 1$. Since $T_i^2 = U_i^2$, $t_i \in U_i^2$ contradicting the requirements of the Borel cover.

Case 2: $T_i^k \neq U_i^k$ for some i and some $k \geq 2$. Choose the player i and number k so that $T_j^\ell = U_j^\ell$ for each player j and each ℓ with $k > \ell \geq 2$. Note, we must either have $T_i^k \backslash U_i^k \neq \emptyset$ or $T_i^k \subsetneq U_i^k$.

First suppose that $T_i^k \backslash U_i^k \neq \emptyset$. If so, we can find some type of some player $j \neq i$, viz. $t_j \in T_j^{k+1}$ so that (i) $\text{marg}_{S_{-j}} \beta_j^{**}(t_j) \neq \mu_j$, (ii) $\beta_j^{**}(t_j)(S_{-j} \times T_{-j-i} \times (T_i^k \backslash U_i^k)) > 0$, and (iii) $\beta_j^{**}(t_j)(S_{-j} \times T_{-j-i} \times U_i^k) > 0$.

Since $\{U_i^\ell : \ell = 1, 2, \ldots\}$ is a partition, there is not U_j^m that can contain t_j and satisfy the properties of a Borel cover.

Next suppose that $T_i^k \subsetneq U_i^k$. Then there exists some $t_i \in U_i^k$ with $\beta_i^{**}(t_i)(S_{-i} \times T_{-i}^{k-1}) = 1$ but $t_i \notin T_i^k$. This contradicts the construction of the covers described in the proof. □

Appendix C Proofs for Section 6

Proof of Proposition 6.1. Begin with part (i). Fix some $s_i \in \mathrm{proj}_{S_i}\left(R_i^1 \cap (S_i \times T_i^1)\right)$. Then there exists some $t_i \in T_i^1$ so that $(s_i, t_i) \in R_i^1$. As such, $s_i \in \mathbb{BR}_i[\mathrm{marg}_{S_{-i}}\beta_i(t_i)]$ and $\mathrm{marg}_{S_{-i}}\beta_i(t_i) = \mu_i$. So $s_i \in L_i^1[\mu]$. Conversely, fix $s_i \in L_i^1[\mu]$. Then $s_i \in \mathbb{BR}_i[\mu_i]$ and, for each $t_i \in T_i^1$, $\mathrm{marg}_{S_{-i}}\beta_i(t_i) = \mu_i$. Thus, $\{s_i\} \times T_i^1 \subseteq R_i^1 \cap (S_i \times T_i^1)$. As such, $L_i^1[\mu] \subseteq \mathrm{proj}_{S_i}\left(R_i^1 \cap (S_i \times T_i^1)\right)$.

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

Proof of Theorem 6.1. The proof is by induction on k. The case of $k = 1$ is part (i) of Proposition 6.1. So, assume the result holds for k. By part (ii) of Proposition 6.1, it suffices to show that

$$L_i^{k+1}[\mu] \subseteq \mathrm{proj}_{S_i}\left(R_i^{k+1} \cap (S_i \times T_i^{k+1})\right).$$

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

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

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

Appendix D Proof for Section 7

Proof of Theorem 7.1. It is standard that, for each m, $\mathrm{proj}_S R^m \subseteq S^m$. To establish the converse, 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.

$m = 1$: 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 $\mathrm{marg}_{S_{-i}}\beta_i(t_i^2) = \nu_i$. As such, $(s_i, t_i^2) \in R_i^1$.

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

Appendix E Proofs for Section 8

E.1 Finite-Order Belief Type Structures

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

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

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

Say $(s_i, \tilde{t}_i)$ is **rational** if $\tilde{\beta}_i(\tilde{t}_i) \in \Delta(S_{-i} \times (\tilde{T}_{-i} \cup \{d\}))$ and satisfies the condition in Definition 3.3. Say $\tilde{t}_i$ **believes** an event E_{-i} if $\tilde{\beta}_i(\tilde{t}_i) \in \Delta(S_{-i} \times (\tilde{T}_{-i} \cup \{d\}))$ and $\tilde{t}_i$ satisfies the condition in Definition 3.4.

We define RmBR analogously to Definition 3.5. In particular, we write $\tilde{R}_i^1$ for the set of rational strategy-type pairs and $\tilde{R}_i^{m+1}$ for the set of strategy-type pairs which satisfy rationality and m^{th}-order belief of rationality.

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

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

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

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

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

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

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

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

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

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

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

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

In fact, we will show a slightly stronger claim:

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

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

The proof is by induction on m.

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

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

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

E.2 Alternate Definitions of a Complete Type Structures

In this section, we show that we cannot replace condition (iv) of Definition 5.2 by condition (iv”). Return to the Undercutting Game and take $\mu = (\mu_1, \mu_2)$ so that $\mu_1(X) = \mu_2(X) = 1$. Note, there are maps

$g_i : S_i \to \Delta(S_{-i})$ so that $s_i \notin \mathbb{BR}_i[g_i(s_i)]$. We will construct a type structure so that there are Borel covers satisfying conditions (i)-(ii)-(iii) and (iv") but for which $\text{proj}_{S_i}(R_i^k \cap (S_i \times T_i^k)) \subsetneq L_i^k[\mu]$ for some k. Since each $L_i^k[\mu]$ is non-empty, it suffices to construct such a type structure where some $R_i^k \cap (S_i \times T_i^k) = \emptyset$.

For each $k \geq 1$, let $(V_i^k : k \geq 1)$ be a collection of uncountable disjoint compact metric spaces. Write $V_i = \bigcup_{k \geq 1} V_i^k$ and note that V_i is a metric space. **[NOTE:** cite] For each $k \geq 1$, define functions $f_i^k : V_i^k \to \Delta(S_{-i})$ as follows: For each $t_i \in V_i^1$, set $f_i^1(t_i)(X) = 1$ and observe that this function is continuous. For each $k \geq 2$, choose f_i^k to be a continuous onto function. The fact that such functions exist follows from **[NOTE:** cite] .

Next we define functions $\tau_{-i}^k : S_{-i} \times S_{-i} \times V_{-i}$ to satisfy the following requirements. First, for each $s_{-i} \in S_{-i}$, $\tau_{-i}^1(s_{-i}) \subseteq \{s_{-i}\} \times V_{-i}^1$. Second, for each $s_{-i} \in S_{-i}$ each $k \geq 2$, $\tau_{-i}^k(s_{-i}) \subseteq \{s_{-i}\} \times V_{-i}^{k-1}$. Third, for each $s_{-i} \in S_{-i}$, $\tau_{-i}^3(s_{-i}) = (s_{-i}, v_{-i}^2)$ where $f_{-i}^2(v_{-i}^2) = g_{-i}(s_{-i})$. (Note, we can find such a v_{-i}^2 since f_{-i}^2 is onto.) Since each $\tau_{-i}^k(S_{-i})$ is a finite set, each of these functions are continuous. Thus, we can define functions $\underline{f}_i^k : V_i^k \to \Delta(S_{-i} \times V_{-i})$ so that each $\underline{f}_i^k(v_i)$ is the image measure of $f_i^k(v_i)$ under τ_{-i}^k. These functions are themselves continuous. (See Aliprantis and Border, 2007, Theorem 15.14.)

Now construct a type structure $\mathcal{T}$ so that the type sets are $T_i = V_i$ and, for each $k \geq 1$ and $t_i \in V_i^k$, $\beta_i(t_i) = \underline{f}_i^k(t_i)$. Note that β_i is measurable: For each measurable $E \subseteq S_{-i} \times T_{-i}$, $(\beta_i)^{-1}(E) = \bigcup_{k \geq 1} (\underline{f}_i^k)^{-1}(E)$ is a countable union of measurable sets and so measurable. Moreover, the collection $\{V_i^k : k = 1, 2, \ldots\}$ is a Borel cover of V_i and the collections jointly satisfy conditions (i)-(ii)-(iii) and (iv").

Finally, observe that $R_i^3 \cap (S_i \times V_i^3) = \emptyset$: Fix $t_i \in V_i^3$ and note that

$$\text{Supp}\, \beta_i(t_i) \subseteq \{(s_{-i}, v_{-i}^2) \in S_{-i} \times V_{-i}^2 : \text{marg}_{S_{-i}} \beta_{-i}(v_{-i}^2) = g_{-i}(s_{-i})\}.$$

Thus, t_i cannot believe R_{-i}^1. As a consequence, $R_i^2 \cap (S_i \times V_i^3)$.

E.3 RmBR Behavior of k-Types

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

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

for each m.

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

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

Proof of Lemma 8.1. By Lemma E.1,

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

from which the claim follows. ■

References

Alaoui, L. and A. Penta. 2014. "Endogenous Depth of Reasoning.". `http://www.econ.wisc.edu/~apenta/research.html`.

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.

Aumann, R.J. 1987. "Correlated equilibrium as an expression of Bayesian rationality." *Econometrica: Journal of the Econometric Society* pp. 1–18.

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

Bernheim, B.D. 1984. "Rationalizable Strategic Behavior." *Econometrica* 52(4):1007–1028.

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.

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

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

Crawford, Vincent P, Miguel A Costa-Gomes and Nagore Iriberri. 2013. "Structural Models of Nonequilibrium Strategic Thinking: Theory, Evidence, and Applications." *Journal of Economic Literature* 51(1):5–62.

Friedenberg, A. and H. Jerome Keisler. forthcoming. "Iterated Dominance Revisited.".

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

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/`.

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.

Pearce, D.G. 1984. "Rationalizable Strategic Behavior and the Problem of Perfection." *Econometrica* 52(4):1029–1050.

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.

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.