How Many Levels Do Players Reason? An Observational Challenge and Solution*

Adam Brandenburger† Alex Danieli‡ Amanda Friedenberg§

May 10, 2017

Abstract

How many levels do players reason about rationality? In simultaneous-move games, there is a well-understood methodology for identifying reasoning about rationality—one based on iterated dominance. However, there is a challenge in porting that methodology to extensive-form games: It would appear to require that the researcher have information about the hierarchies of beliefs that players consider possible. Moreover, imposing (potentially incorrect) assumptions about those beliefs can have non-trivial implications for identification—implications that do not arise in simultaneous-move games. We provide a novel methodology to identify reasoning in a generic class of extensive-form games. Importantly, the methodology does not require information about the hierarchies of beliefs players consider possible. The centipede game illustrates that the methodology has non-trivial implications for experimental design.

1 Introduction

Interactive reasoning is an important aspect of behavior. To determine whether a particular course of action is good or bad, a player Ann may need to form a belief about her co-player Bob's behavior. In forming such a belief, she may reason that Bob is "strategically sophisticated." For instance, she may form her belief about Bob's behavior by reasoning that he is "rational," i.e., that he maximizes his expected utility given his belief about Ann's behavior. Or, she may form her belief by reasoning that Bob is "rational and reasons about rationality." That is, she may form her belief by reasoning that Bob is rational and that he, in turn, forms his belief (about Ann's behavior) by reasoning that she is rational. And so on.

A natural theoretical benchmark is that players are rational and their beliefs (about behavior) are consistent with "common belief of rationality." Under this benchmark, each player believes that their co-player is rational, that their co-player believes that they are rational, and so on, ad infinitum. More loosely, this benchmark involves a player reasoning, at all levels, about their co-player's rationality. But,

*We thank Vince Crawford, Jessy Hsieh, Terri Kneeland, Xiaomin Li, and Elliot Lipnowski for valuable comments. We also thank audiences at Bristol, Oxford, Penn State, Washington University in St. Louis, McGill, National University of Singapore, the 2013 Asian Meeting of the Econometric Society, the 2013 Workshop on Logical Structure of Correlated Information Change, SAET, TARK, and the Game Theory Society World Congress for important input. Adam Brandenburger thanks the NYU Stern School of Business, NYU Shanghai, and J.P. Valles for financial support. Amanda Friedenberg and Alex Danieli thank NSF SES-1358008 for financial support. Amanda Friedenberg thanks the UCL Economics Department for its unbounded hospitality.

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

‡W.P. Carey School of Business, Arizona State University, Alexander.Danieli@asu.edu

§W.P. Carey School of Business, Arizona State University, amanda.friedenberg@asu.edu, amandafriedenberg.org

in practice, players may depart from such a benchmark. For instance, Ann and Bob may have previously interacted—either directly or with a population of likeminded players. In the course of that interaction, Ann may have observed behavior that she could not rationalize. If so, she may, instead, form her belief by reasoning that Bob is irrational. Or, she may form her belief by reasoning that "Bob is rational but believes she is irrational." And so on.

This raises the question: How many levels do players reason about rationality? Addressing this question is fundamental to the analysis of games. In particular, the answer can serve to generate new—i.e., out-of-sample—predictions. This, in turn, can have important normative implications for the design of markets, institutions, and policies.

To address this question, we focus on the case where the researcher only observes behavior of the players—or, perhaps, only observes the outcome of the game. We provide a methodology to identify reasoning in a broad class of extensive-form games.

Benchmark: Simultaneous Move Games The methodology for identifying reasoning about rationality is well understood in simultaneous move games. As such, it will be useful to begin with the simultaneous-move benchmark.

In simultaneous move games, reasoning about rationality is typically formulated as *rationality and m^{th}-order belief of rationality*: A player is *rational* if she chooses a best response, given her belief about how the game is played. A player is rational and 1^{st}-order *believes* rationality, if she is rational and has a belief that assigns probability one to her co-players' rationality. Inductively, a player is rational and m^{th}-order believes rationality if she is rational and n^{th}-order believes rationality, for all $n = 1, \ldots, m-1$.

An epistemic framework can be used to formally capture rationality and m^{th}-order belief of rationality. (See, e.g., Brandenburger, 2007 and Dekel, Siniscalchi et al., 2014 for an overview.) The associated behavior is captured by the $(m+1)$-*rationalizable* strategies (Bernheim, 1984; Pearce, 1984). Specifically:

> **Theorem**: A strategy is consistent with rationality and m^{th}-order belief of rationality if and only if the strategy is $(m+1)$-rationalizable. A strategy is consistent with rationality and common belief of rationality if and only if the strategy is rationalizable.

(See, e.g., Tan and Werlang, 1988.) This standard result can be used to provide a bound on reasoning about rationality: If the researcher observes a subject choose an m- but not $(m+1)$-rationalizable strategy, the researcher can conclude that the subject's behavior is consistent with rationality, $(m-1)^{\text{th}}$-order belief of rationality, but not m^{th}-order belief of rationality. Likewise, if the researcher observes a subject choose a rationalizable strategy, the researcher can conclude that her behavior is consistent with rationality and common belief of rationality.

To sum up, in simultaneous games, we can use observed behavior to infer the maximum level of reasoning (about rationality) consistent with observed behavior. In light of well-known results, this can be achieved by inferring the maximum m so that the behavior is m-rationalizable.

Beyond Simultaneous Move Games The goal is to implement the previous methodology for a broad class of extensive-form games. To do so, it is useful to highlight the structure behind the simultaneous-move game analysis: First, we specify an epistemic framework. Within the framework, we formulate what it means to reason about rationality. This corresponded to rationality and m^{th}-order belief of rationality (or rationality and common belief about rationality). Finally, as a theorem, we derive behavioral implications of such reasoning. This came in the form of the $(m + 1)$-rationalizable strategies.

Strictly speaking, the epistemic framework is a prerequisite for specifying rationality and m^{th}-order belief of rationality. (This point will be highlighted in Section 3.) That said, one might wonder if it is really necessary from the perspective of identifying reasoning about rationality. For instance, in the context of simultaneous move games, one might intuite that m-rationalizability can be used to bound the level of reasoning about rationality—that is, one would have intuitively come to that conclusion, even if one did not take the steps of going through an epistemic framework.

Likewise, one might conjecture that, if the extensive-form game is sufficiently "simple," then such an epistemic approach is not needed. For instance, one might conjecture that, in the context of perfect-information games, one can analogously skip the step of going through an epistemic framework: that the backward induction algorithm can be used to provide such a bound. However, this idea cannot be done quite generally—at least not in an obvious way.

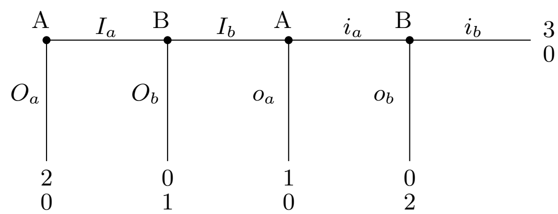

Figure 1.1: Reny (1992)

Consider the game in Figure 1.1. Observe that the strategy I_a-o_a is dominated at the beginning of the game by O_a. As such, there is no belief that A can hold, so that I_a-o_a would be a best response for A. So, if we observe a subject in the role of A play I_a-o_a, we should conclude that this subject is irrational. That said, the backward induction algorithm cannot easily be used provide such a bound. No strategy of A is eliminated on the first round of the backward induction algorithm. Moreover, the second round of the backward induction algorithm eliminates I_a-i_a and not I_a-o_a. In fact, I_a-o_a is not eliminated until the fourth and final round of the algorithm.

In Figure 1.1, we could not use backward induction to provide a bound on the level of reasoning. This arose from a fundamental feature about backward induction: The algorithm does not eliminate all dominated (i.e., all sequentially irrational) strategies on the first round of the procedure. (Note, then, this conclusion does not rely on details of epistemic assumptions, i.e., details about how "reasoning about rationality" is formulated.) Later we will also show that extensive-form rationalizability Pearce (1984) cannot be used to provide such a bound. (See Examples 3.2-3.3.) That point will be more subtle and, in particular, will depend on reasoning about rationality.

This Paper We will we specify an epistemic framework appropriate for studying extensive-form games. This allows us to formalize how players reason about rationality. We focus on *rationality and m^{th}-order strong belief of rationality* (Battigalli and Siniscalchi, 2002). (We discuss this concept and choice in Sections 3 and 9.) Theorem 6.1 provides a behavioral characterization of rationality and m^{th}-order strong belief of rationality that applies to any generic game. In particular, it shows that, in such games, the behavior is characterized by, what we will call, the *m-best response property*. Sections 4 and 8 use this result to provide an algorithm for identifying the maximum the level of reasoning about rationality consistent with observed behavior.

Section 7 applies the results to the Centipede game and, in so doing, shows that there are non-trivial implications for experimental design. We show that, by having a subject play a series of Centipede games in the role of the first player, the researcher can disentangle any m from $(m+1)^{th}$-order reasoning. However, importantly, the researcher cannot use behavior in the role of the second-player to make non-trivial inferences about reasoning. At least, the researcher cannot do so absent imposing auxiliary assumptions about the players' beliefs.

Literature There is a long tradition of using behavior to infer levels of iterative reasoning. Specifically, a prominent literature seeks to infer such iterative reasoning by using level-k and cognitive hierarchy models (e.g. Nagel, 1995; Stahl and Wilson, 1995; Costa-Gomes, Crawford and Broseta, 2001; Camerer, Ho and Chong, 2004; Costa-Gomes and Crawford, 2006). A central feature of their identification strategy is that they impose auxiliary assumptions about beliefs.[1] (This comes in the form of assumptions about the behavior of Level-0 types, which pins down the beliefs of Level-1 types.) With this, their notion of iterative reasoning can be conceptually distinct (in subtle ways) from rationality and m^{th}-order belief of rationality (i.e., even in simultaneous move games).[2] However, often, there is a close relationship in terms of observed behavior. (Appendix A in Friedenberg, Kets and Kneeland (2016) discusses this further.)

This paper differs from that literature, in that we do not impose auxiliary assumptions about players beliefs. Examples 3.2-3.3 will highlight the fact that this is important in the context of dynamic games—i.e., in a sense that does not arise in simultaneous move games. Thus, we seek to understand the extent to which the researcher can identify levels of reasoning without imposing such auxiliary assumptions. In certain games, the researcher will not be able to provide a meaningful identification of iterative reasoning absent auxiliary assumptions. However, importantly, in other cases the researcher will be able to identify iterative reasoning. See Section 7.

With this in mind, our approach is closer in spirit to Kneeland (2015). Kneeland identifies reasoning about rationality without imposing auxiliary assumptions about beliefs. So, in Kneeland, iterative reasoning corresponds to rationality and m^{th}-order belief of rationality. However, the two papers have distinct goals: Whereas Kneeland seeks to identify levels of reasoning in a specific simultaneous-move game experiment, this paper seeks to provide a methodology that would allow the researcher to design analogous extensive-form game experiments.

This paper fits in a growing literature on inferring reasoning in dynamic games. For instance, Siniscalchi (2016) and Healy (2015) seek to elicit players beliefs in extensive-form games. (Siniscalchi seeks to elicit conditional hierarchies of beliefs about the strategies and Healy seeks to elicit beliefs about rationality, etc.) By contrast, here, we assume that the researcher does not have access to information about beliefs, only access to data on the play of the game.

2 Epistemic Games

Consider the game in Figure 2.1, Battle of the Sexes (BoS) with an Outside Option. Suppose that we observe Ann play *Out.* We will want to identify the maximum level reasoning about rationality consistent

[1] Importantly, some papers provide evidence in favor of those assumptions, based on auxiliary data. (See, e.g., Costa-Gomes, Crawford and Broseta, 2001, Costa-Gomes and Crawford, 2006, Rubinstein, 2007.)

[2] This arises because Level-0 types may play irrational strategies with positive probability. If so, Level-1 types correspond to a player that is rational but does not believe rationality.

with the observed behavior.

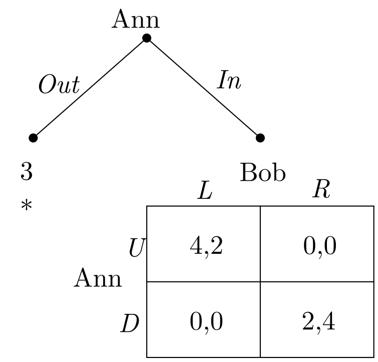

Figure 2.1: Battle of the Sexes with an Outside Option

A starting point will be to formally define what we mean by reasoning about rationality. This will involve describing the strategic situation by, what is called, an epistemic game. To understand why, observe that the strategy *Out* is rational—i.e., a best response—for Ann, if she believes that Bob will play R; it is irrational for her if she believes that Bob will play L. Thus, we cannot specify whether *Out* is rational or irrational without specifying Ann's beliefs about Bob's play of the game.

Consider the case where Ann believes that Bob plays R. This strategy will be rational for Bob if, conditional upon BoS being played, Bob believes that Ann plays *In*-D; but this same strategy is irrational for Bob if, conditional upon BoS being played, Bob believes that Ann plays *In*-U. Thus, to specify whether Ann is "rational and reasons that Bob is rational" we not only need to specify Ann's belief about the strategy Bob plays but also her belief about Bob's belief about her own play.

Continuing along these lines, we need to specify Ann's hierarchies of beliefs about the play of the game. In what follows, we will describe these hierarchies of beliefs by way of a type structure. This is what we do below.

2.1 Extensive-Form Game

Write Γ for a finite two-player extensive-form game of perfect recall, in the sense of Kuhn (1953). The players of the game are a (Ann) and b (Bob).[3] Write c for an arbitrary player in $\{a, b\}$ and $-c$ for the player in $\{a, b\} \backslash \{c\}$. The underlying game tree has an initial node of ϕ and a set of terminal nodes Z. Write H_c for the set of information sets at which player c moves. Assume the game is **non-trivial**, in the sense that each player has at least two distinct actions at some $h \in H_c$. The set of information sets, viz. $H = H_a \cup H_b$, forms a partition of the non-terminal nodes. Player c's extensive-form payoff function is given by $\Pi_c : Z \to \mathbb{R}$.

Let S_c be the set of strategies for player c and let $S = S_a \times S_b$. There is a mapping $\zeta : S \to Z$ so that $\zeta(s_a, s_b)$ is the terminal node reached by (s_a, s_b). Player c's strategic-form payoff function is given by $\pi_c = \Pi_c \cdot \zeta$.

Say a strategy profile (s_a, s_b) **reaches** $h \in H$ if the path from ϕ to $\zeta(s_c, s_{-c})$ passes through some node in h. Write $S(h)$ for the set of strategy profiles that reach h and write $S_c(h) = \text{proj}_{S_c} S(h)$. If a strategy

[3]The analysis extends to three or more players, up to issues of correlation. See Section 9f.

$s_c \in S_c(h)$, then we say that s_c **allows** $h \in H$.

In what follows it will be convenient to look at product sets. Say $Q \subseteq S$ is a **product set** if $Q = \text{proj}_{S_a} Q \times \text{proj}_{S_b} Q$. We take the convention that, if $Q = \emptyset$, then both $\text{proj}_{S_a} Q = \text{proj}_{S_b} Q = \emptyset$.

2.2 Type Structure

As discussed above, a type structure will be used to model strategic uncertainty, i.e., uncertainty about the play of the game. In a type structure, each player will have hierarchies of beliefs about the play of the game. Because the focus is on extensive-form games, over the course of playing the game, players may learn information inconsistent with their initial hypothesis. For an example of this, return to BoS with the Outside Option: Bob may begin the game with a hypothesis that Ann exercises her outside option, but may come to learn that this is false. If so, he will be forced to revise his belief about Ann's strategy choice. In light of this, we will need to specify conditional beliefs about the play of the game.

Refer to $(\Omega, \mathcal{B}(\Omega))$ as a probability space, when Ω is a compact metric space and $\mathcal{B}(\Omega)$ is the Borel sigma-algebra on Ω. Write $\mathcal{P}(\Omega)$ for the set of Borel probability measures on Ω. Endow $\mathcal{P}(\Omega)$ with the topology of weak convergence so that it is again a compact metric space.

Call $(\Omega, \mathcal{B}(\Omega), \mathcal{E})$ a **conditional probability space** if $(\Omega, \mathcal{B}(\Omega))$ is a probability space and $\mathcal{E} \subseteq \mathcal{B}(\Omega) \backslash \{\emptyset\}$ is finite. The collection $\mathcal{E}$ will be referred to as (a finite set of) **conditioning events**. Since $\mathcal{B}(\Omega)$ is clear from the context, we will suppress reference to $\mathcal{B}(\Omega)$ and simply write $(\Omega, \mathcal{E})$ for a conditional probability space.

Definition 2.1. Fix a conditional probability space $(\Omega, \mathcal{E})$. An **array** on $(\Omega, \mathcal{E})$ is some $p : \mathcal{B}(\Omega) \times \mathcal{E} \to [0, 1]$ so that, for each $E \in \mathcal{E}$, $p(\cdot|E) \in \mathcal{P}(\Omega)$ with $p(E|E) = 1$.

Definition 2.2. Fix a conditional probability space $(\Omega, \mathcal{E})$. A **conditional probability system** (**CPS**) on $(\Omega, \mathcal{E})$ is an array $p : \mathcal{B}(\Omega) \times \mathcal{E} \to [0, 1]$ that satisfies the following criterion: If $E, F \in \mathcal{E}$ with $G \subseteq F \subseteq E$, then $p(G|E) = p(G|F)p(F|E)$.

Write $\mathcal{A}(\Omega, \mathcal{E})$ for the set of arrays on $(\Omega, \mathcal{E})$ and write $\mathcal{C}(\Omega, \mathcal{E})$ for the set of CPS's on $(\Omega, \mathcal{E})$. Note that $\mathcal{C}(\Omega, \mathcal{E}) \subseteq \mathcal{A}(\Omega, \mathcal{E}) \subseteq [\mathcal{P}(\Omega)]^{|\mathcal{E}|}$. Endow $[\mathcal{P}(\Omega)]^{|\mathcal{E}|}$ with the product topology and $\mathcal{C}(\Omega, \mathcal{E})$ with the relative topology, so that $\mathcal{C}(\Omega, \mathcal{E})$ is a compact metric space.

In our analysis, player c's set of conditioning events will correspond to

$$\mathcal{E}_c = \{S_{-c}(h) : h \in H_c \cup \{\phi\}\}.$$

So, Ann has a conditioning event that corresponds to the beginning of the game, viz. $S_b(\phi) = S_b$. Moreover, she also has conditioning events $S_b(h)$ corresponding to each information set $h \in H_a$ at which she moves.

Definition 2.3. A **Γ-based type structure** is some $\mathcal{T} = (\Gamma; T_a, T_b; \beta_a, \beta_b)$ where

(1) T_c is a compact metric **type set** for player c and

(2) $\beta_c : T_c \to \mathcal{C}(S_{-c} \times T_{-c}, \mathcal{E}_c \otimes T_{-c})$ is a continuous **belief map** for player c.

A Γ-based type structure models hierarchies of conditional beliefs about the play of the game: A type of Ann, viz. t_a, is associated with a CPS $\beta_a(t_a)$ on $(S_b \times T_b, \mathcal{E}_a \otimes T_b)$. As such, the type it also associated with a CPS about Bob's play, i.e., a first-order CPS on $(S_b, \mathcal{E}_a)$. (See Lemma A.1.) Since each type of Bob is also associated with a first-order CPS on $(S_a, \mathcal{E}_b)$, type t_a of Ann is associated with a second-order CPS, i.e., a system of beliefs about both (i) Bob's play S_b and (ii) Bob's first-order CPS's on $(S_a, \mathcal{E}_b)$. And so on.

For any given game Γ, there are many Γ-based type structures. Write $\mathbb{T}(\Gamma)$ for the class of Γ-based type structures. Battigalli and Siniscalchi (1999) construct a canonical type structure, which induces all possible hierarchies of conditional beliefs. Their type structure $\mathcal{T}^* = (\Gamma; T_a^*, T_b^*; \beta_a^*, \beta_b^*)$ has the property that it is **type-complete** (Brandenburger, 2003), i.e., for each CPS $\mathrm{p}_c \in \mathcal{C}(S_{-c} \times T_{-c}^*, \mathcal{E}_c \otimes T_{-c}^*)$, there is a type t_c with $\beta_c(t_c) = \mathrm{p}_c$. Other type structures model an assumption that some event is common (full) belief. (See Appendix A in Battigalli and Friedenberg, 2009 for a formal treatment.) The following gives an example.

Example 2.1. Consider BoS with an Outside Option. Suppose that it is commonly understood that "Bob is a bully" and, so, whenever a BoS game is played, he attempts to go for his best option and play R.[4] In particular:

Bully-1: at each information set, Ann believes that Bob plays R,

Bully-2: at each information set, Bob believes "Bully-1,"

Bully-3: at each information set, Ann believes "Bully-2,"

etc. This is a restriction on the hierarchies of beliefs that the players consider possible.

This restriction on the hierarchies of beliefs can be captured by a type structure $\mathcal{T} = (\Gamma; T_a, T_b; \beta_a, \beta_b)$ satisfying the following properties:

- Each $\beta_a(t_a)(\cdot|S_b \times T_b)$ assigns probability one to $\{R\} \times T_b$.
- For each CPS p_a with $p_a(\{R\} \times T_b | S_b \times T_b) = 1$, there is a type t_a with $\beta_a(t_a) = \mathrm{p}_a$.
- For each CPS p_b, there is a type t_b with $\beta_b(t_b) = \mathrm{p}_b$.

The fact that such a type structure exists follows from Battigalli and Friedenberg (2009). It captures the contextual assumption that there is common (full) belief that Bob plays $\{R\}$. □

Remark 2.1. As is standard, we identify a simultaneous-move game with an extensive-form game in which all players move without information about past play. In that case, $\mathcal{E}_a = \{S_b\}$ and $\mathcal{E}_b = \{S_a\}$. Thus, $\mathcal{C}(S_{-c} \times T_{-c}, \mathcal{E}_c \otimes T_{-c}) = \mathcal{P}(S_{-c})$, i.e., $\beta_c : T_c \to \mathcal{P}(S_{-c} \times T_{-c})$. □

2.3 Epistemic Game

For a given game Γ, write $\mathbb{T}(\Gamma)$ for the set of Γ-based type structures. Since Γ is non-trivial, there is an uncountable number of elements in $\mathbb{T}(\Gamma)$. An (**extensive-form**) **epistemic game** is some pair $(\Gamma, \mathcal{T})$ with $\mathcal{T} \in \mathbb{T}(\Gamma)$. The epistemic game is the exogenous description of the strategic situation. An epistemic game induces a set of **states**, viz. $S_a \times T_a \times S_b \times T_b$.

In what follows, we will fix an extensive-form game Γ. With this, the epistemic game can be identified with a type structure in $\mathbb{T}(\Gamma)$. As such, we often conflate 'type structure' with 'epistemic game.' No confusion should result.

[4] We use the phrase "common understanding" for "common full belief." Example 3.1 defines the concept of "full belief."

3 Epistemic Conditions

The set of states associated with a type structure $\mathcal{T}$ describes a set of possible strategy-type pairs that can obtain. This alone imposes no restrictions on behavior or strategic reasoning. The epistemic conditions will do just that. We begin by imposing the behavioral conditional of rationality. We then impose strategic reasoning.

3.1 Rationality

Fix some $X_c \subseteq S_c$. Say $s_c \in X_c$ is a **best response under** $\mu \in \mathcal{P}(S_{-c})$ **given** X_c if

$$\sum_{s_{-c} \in S_{-c}} [\pi_c(s_c, s_{-c}) - \pi_c(r_c, s_{-c})]\mu(s_{-c}) \geq 0$$

for all $r_c \in X_c$.

Definition 3.1. Say s_c is a **sequential best response under** $\mathrm{p}_c \in \mathcal{A}(S_{-c}, \mathcal{E}_c)$ if, for each $h \in H_c$ with $s_c \in S_c(h)$, s_c is a best response under $p_c(\cdot|S_{-c}(h))$ given $S_c(h)$.

Notice that each $\beta_c(t_c)$ induces a CPS in $\mathcal{C}(S_{-c}, \mathcal{E}_c)$ via marginalization. Specifically, the marginal CPS $\text{marg}_{S_{-c}} \beta_c(t_{-c}) \in \mathcal{C}(S_{-c}, \mathcal{E}_c)$ is a CPS p_c with $p_c(\cdot|S_{-c}(h)) = \text{marg}_{S_{-c}} \beta_c(t_c)(\cdot|S_{-c}(h) \times T_{-c})$ for each $S_{-c}(h) \in \mathcal{E}_c$.

Definition 3.2. Say (s_c, t_c) is **rational** if s_c is a sequential best response under the marginal CPS $\text{marg}_{S_{-c}} \beta_c(t_c)$.

3.2 Reasoning about Rationality

We next impose the requirement that a player "reasons" that the other player is rational. We take "reasons" to mean that a player maintains a hypothesis that the other player is rational, so long as she has not observed evidence that contradicts rationality. This idea is captured by strong belief of rationality.

Definition 3.3 (Battigalli and Siniscalchi, 2002)**.** Say a CPS $\mathrm{p} \in \mathcal{C}(\Omega, \mathcal{E})$ **strongly believes** an event F if, for each conditioning event $E \in \mathcal{E}$, $E \cap F \neq \emptyset$ implies $p(F|E) = 1$.

Definition 3.4. A type t_c **strongly believes** an event $E_{-c} \subseteq S_{-c} \times T_{-c}$ if $\beta_c(t_c)$ strongly believes E_{-c}.

Notice, in the specific case of a simultaneous-move game, "strong belief" coincides with "belief," i.e., a type believes an event E_{-c} if its single probability measure assigns probability one to the event E_{-c}.

To better understand the concept of strong belief, it will be useful to contrast it with "full belief." The next example describes full belief and why we focus, instead, on strong belief.[5]

Example 3.1. Consider the game in Figure 3.1. Observe that, in this game, Ann's choice of I_a is dominated. Thus, for any associated epistemic game and any type t_a, (s_a, t_a) is rational if and only if $s_a = O_a$. Similarly, for any epistemic game and any type t_b, (s_b, t_b) is rational if and only if $s_b = O_b$.

Say a type of Bob t_b *fully believes* an event $E_a \subseteq S_a \times T_a$ if $\beta_b(t_b)$ assigns probability one to E_a given every conditioning event. For any type structure, there is no type of Bob that fully believes that Ann is

[5]See Section 9b for a discussion of why we focus on strong belief instead of the alternate concept of initial belief.

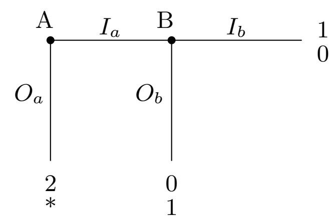

Figure 3.1: Strong Belief versus Full Belief

rational. This is because, conditional upon Bob reaching the information set associated with Ann's choice of I_a, Bob must assign probability one to the event that Ann is irrational.

Suppose we were to take reasoning about rationality as "full belief of rationality." In that case, when we observe Bob play O_b, we would conclude that Bob's behavior is consistent with rationality, but not not reasoning about rationality. Thus, we would view Bob's rationality bound as 1. However, this is an artifact of the impossibility of believing Ann is rational, conditional upon observing Ann play an irrational move. It does not reflect a lack of reasoning on Bob's part.

Instead, consider the following type structure: There is one type for each player, i.e., $T_a = \{t_a\}$ and $T_b = \{t_b\}$. The belief map of Ann specifies $\beta_a(t_a)((O_b, t_b)|S_b \times T_b) = 1$. The belief map of Bob specifies $\beta_b(t_b)((O_a, t_a)|S_a \times T_a) = 1$ and $\beta_b(t_b)((I_a, t_a)|\{I_a\} \times T_a) = 1$. In this case, (O_a, t_a) and (O_b, t_b) are both rational, as they must be. Moreover, t_a strongly believes "Bob is rational" and, importantly, now t_b also strongly believes "Ann is rational." (That is, t_b maintains this hypothesis, so long as it is not contradicted by Ann's behavior.) Inductively, we can see that at (O_a, t_a, O_b, t_b) there is "rationality and common strong belief of rationality." (The concept will be defined shortly.) □

3.3 Rationality and m^{th}-order Strong Belief of Rationality

Fix some type structure $\mathcal{T} \in \mathbb{T}(\Gamma)$. Set $R_c^0(\mathcal{T}) = S_c \times T_c$. Let $R_c^1(\mathcal{T})$ be the set of rational strategy-type pairs (s_c, t_c) in $\mathcal{T}$. Inductively define sets $R_a^m(\mathcal{T})$ and $R_b^m(\mathcal{T})$ so that

$$R_c^{m+1}(\mathcal{T}) = R_c^m(\mathcal{T}) \cap [S_c \times \{t_c : t_c \text{ strongly believes } R_{-c}^m(\mathcal{T})\}].$$

Write $R^m(\mathcal{T}) = R_a^m(\mathcal{T}) \times R_b^m(\mathcal{T})$. The set $R^{m+1}(\mathcal{T})$ is the set of strategy-type pairs (in $\mathcal{T}$) at which there is **rationality and m^{th}-order strong belief of rationality** (**RmSBR**). The set $R^\infty(\mathcal{T}) = \cap_{m \in \mathbb{N}} R^m(\mathcal{T})$ is the set of strategy-type pairs (in $\mathcal{T}$) at which there is **rationality and common strong belief of rationality** (**RCSBR**).

Observe that $\text{proj}_{S_a \times S_b} R^{m+1}(\mathcal{T})$ is the **set of RmSBR predictions** for the type structure $\mathcal{T}$. One natural conjecture is that this gives rise to the set of extensive-form rationalizable strategies (EFR, Pearce, 1984). Extensive-form rationalizability sequentially eliminates strategies that are not sequential best responses. Battigalli and Siniscalchi (2002, Proposition 6) show that, when the type structure is $\mathcal{T}$ is type-complete, the set of RmSBR predictions is the set of strategies that survive $(m+1)$ rounds of extensive-form rationalizability. However, this need not be the case if $\mathcal{T}$ is type-incomplete. In that case, the predictions of RmSBR may be disjoint from the extensive-form rationalizable (EFR) strategies. The following examples illustrate this point.

Example 3.2. Consider BoS with an Outside Option. There is a CPS on S_a so that L (resp. R) is a sequentially best response—i.e., a CPS that assigns probability one to *In*-U (resp. *In*-D) conditional upon *In* being played. Likewise, there is a CPS on S_b so that *Out* (resp. *In*-U) is a sequential best response—i.e., a CPS that assigns probability one to R (resp. L). But there is no CPS so that *In*-D is a sequential best response. Thus, one round of EFR gives

$$\mathrm{EFR}_a^1 \times \mathrm{EFR}_b^1 = \{Out, In\text{-}U\} \times \{L, R\}.$$

Now observe that there is no CPS on EFR_a^1 so that R is a sequential best response. Thus,

$$\mathrm{EFR}_a^2 \times \mathrm{EFR}_b^2 = \{Out, In\text{-}U\} \times \{L\}.$$

With this,

$$\mathrm{EFR}_a^3 \times \mathrm{EFR}_b^3 = \{In\text{-}U\} \times \{L\}.$$

Thus, there is one EFR strategy profile, (*In*-U,L).

Battigalli and Siniscalchi (2002) show that, if $\mathcal{T}^*$ is type-complete, then EFR corresponds round-for-round to RmSBR in the associated epistemic game. That is, for each m,

$$\mathrm{proj}_{S_a \times S_b} R^m(\mathcal{T}^*) = \mathrm{EFR}_a^m \times \mathrm{EFR}_b^m.$$

As such, the EFR predictions above are also RmSBR predictions, provided the epistemic game is type complete. □

Example 3.3. Again, consider BoS with an Outside Option. Let $\mathcal{T}$ be the type structure from Example 2.1, representing the case where it is commonly understood that "Bob is a bully." Now

$$\mathrm{proj}_{S_a \times S_b} R^m(\mathcal{T}) = \{Out\} \times \{L, R\}$$

for each $m \geq 1$. So, for each $m \geq 1$, there are types t_a^m and t_b^m so that $(Out, t_a^m, R, t_b^m) \in R^m(\mathcal{T})$.

To understand why, observe that (s_a, t_a) is rational if and only if $s_a = Out$. Thus, $R_a^1(\mathcal{T}) = \{Out\} \times T_a$. Now, observe that there is a type $t_b^2 \in T_b$ that, at the initial node, assigns probability one to $\{Out\} \times T_a$ and, conditional upon Ann playing *In*, assigns probability one to $\{In\text{-}D\} \times T_a$. Certainly (R, t_b^2) is rational. In addition, t_b^2 strongly believes the event that "Ann is rational": At the initial node, the type assigns probability 1 to the event that "Ann is rational;" moreover, this event is inconsistent with Ann playing *In*. Thus, $(R, t_b^2) \in R_b^2(\mathcal{T})$. Now observe that there is a type $t_a^3 \in T_a$ that assigns probability one to (R, t_b^2) at each information set. With this, $(Out, t_a^3) \in R_a^3(\mathcal{T})$. And so on. □

4 Identifying Level-m Reasoning

The description of the game consists of both the game itself and the type structure. In the ideal case, the researcher would observe the game, the type structure, and the actual state. With this, he could uniquely deduce the level of reasoning consistent with observed behavior. For instance, if the researcher knew that the epistemic game is characterized by $\mathcal{T}$ and the true state is

$$(s_a, t_a, s_b, t_b) \in (R_a^2(\mathcal{T}) \backslash R_a^3(\mathcal{T})) \times (R_b^1(\mathcal{T}) \backslash R_b^2(\mathcal{T})),$$

then the researcher would conclude that Ann reasons (exactly) 2 levels about rationality and Bob reasons (exactly) 1 level about rationality.

However, we are concerned with the case where the researcher only observes the game Γ and the strategy played—or perhaps a signal of the strategy played (e.g., the outcome of the game). As such, the researcher will be concerned with identifying the maximum level of reasoning consistent with observed behavior.

To understand what this involves, return to BoS with an Outside Option. Suppose the researcher observes Ann play *Out*, but not Ann's type. If the researcher knew that the players' type structure were type-complete, then the researcher would infer that the maximum level of reasoning consistent with Ann's behavior is 2. However, if the researcher knew that it was commonly understood that "Bob is a Bully," the researcher would instead infer that Ann's behavior is consistent with common reasoning.

The issue is that the maximum level of reasoning consistent with Ann's behavior depends on the type structure. But, the researcher faces a challenge of not having information about the players' actual type structure. Thus, we will look at the maximum level of reasoning consistent with Ann's behavior, across all type structures. So, if we observe Ann play *Out*, we will identify Ann's behavior as consistent with common reasoning.

4.1 Approach to Identification

Consider a dataset that is indicative of players' choices in a game Γ. Write $\mathbb{D}$ for the set of possible realizations in this dataset and write $\delta_c : S_a \times S_b \to \mathbb{D}$ for the mapping from strategies of c to data realizations. For instance, if the dataset is obtained by implementing the strategy method in a lab experiment, the dataset would consist of the set of reported strategies for player c and the mapping δ_c would be the projection of $S_a \times S_b$ onto S_c. Alternatively, the dataset might consist of information about the path of play (resp. outcome) in the game—e.g., the researcher may observe a path of bids in an auction (resp. who won the auction and at what price). In that case, $\mathbb{D}$ is the set of terminal nodes (resp. outcomes) and δ_c maps strategy profiles to their induced paths of play (resp. induced outcomes).

The researcher will use the mapping δ_c to back out the maximum level of reasoning consistent with the data, from the maximum level of reasoning consistent with a given strategy. With this in mind, say a strategy s_c is **consistent with** m **levels of reasoning** (resp. **common reasoning**) **for** c if there exists some type structure $\mathcal{T}^*$ so that $s_c \in \text{proj}_{S_c} R_c^m(\mathcal{T}^*)$ (resp. $s_c \in \text{proj}_{S_c} \bigcap_m R_c^m(\mathcal{T}^*)$). Say the observed data d is **consistent with** m **levels of reasoning** (resp. **common reasoning**) **for** c if there exists some strategy profile (s_a^*, s_b^*) so that $\delta_c(s_a^*, s_b^*) = d$ and s_c^* is consistent with m levels of reasoning (resp. common reasoning) for c.

Definition 4.1. The observed data $d \in \mathbb{D}$ identifies player c as a **Level-**m **Reasoner** if d is consistent with m levels of reasoning for c but inconsistent with $(m+1)$ levels of reasoning for c. The observed data $d \in \mathbb{D}$ identifies player c as a **Level-**∞ **Reasoner** if, for each m, the observed data $d \in \mathbb{D}$ is consistent with m levels of reasoning for c.

Player c is identified as a Level-m Reasoner if (i) there is some strategy s_c^* consistent with both the data and R(m-1)SBR, but (ii) there is no strategy s_c consistent with both the data and RmSBR. Player c is identified as a Level-∞ Reasoner if, for each m, there is some strategy s_c consistent with both the data and RmSBR. That is, c is identified as a Level-∞ Reasoner if there is no bound on the levels of reasoning consistent with the data.

Remark 4.1. Observe an important subtlety in the definitions: If the data is consistent with common reasoning, then the data identifies the subject as being a Level-∞ Reasoner. However, the converse is not obvious: A strategy s_c may be consistent with RmSBR for all m (i.e., for each m, there may be some type structure $\mathcal{T}^m$ and associated $(s_c, t_c^m) \in R_c^m(\mathcal{T}^m)$). But, this does not immediately imply that s_c is consistent with RCBR (i.e., there exists some epistemic game $\mathcal{T}$ and some $(s_c, t_c) \in \bigcap_m R_c^m(\mathcal{T})$).

Despite the above, in Section 9c, we will show that the converse also holds: If the data identifies the subject as being a Level-∞ Reasoner, then the data must be consistent with common reasoning. That is, if there is no bound on the number of levels of reasoning consistent with the data, then the data must, in fact, be consistent with RCSBR. □

4.2 A Key Step for Identification

The goal is to construct a partition $\mathcal{L}_c^* = \{L_c^{*,0}, L_c^{*,1}, \ldots, L_c^{*,m}, \ldots, L_c^{*,\infty}\}$ on the data, so that $d \in L_c^{*,m}$ if and only if the data reflects a Level-m Reasoner. To do so, it suffices to construct an auxiliary partition

$$\mathcal{L}_c = \{L_c^0, L_c^1, \ldots, L_c^m, \ldots, L_c^\infty\}$$

on the strategies of player c, viz. S_c, so that the following holds:

(i) For each finite m, $s_c \in L_c^m$ if and only if s_c is consistent with m levels of reasoning but not $(m+1)$ levels of reasoning, and

(ii) $s_c \in L_c^\infty$ if and only if s_c is consistent with common reasoning.

Thus, we focus on constructing the partition $\mathcal{L}_c$.[6]

There is a natural approach to constructing the partition $\mathcal{L}_c$. Refer to Figure 4.1. For each Γ-based type structure $\mathcal{T}$, define $S^m(\mathcal{T})$ (resp. $S_c^m(\mathcal{T})$) to be the projection of the strategy-type pairs in $R^m(\mathcal{T})$ onto S (resp. S_c). (Observe that $S^m(\mathcal{T}) = S_a^m(\mathcal{T}) \times S_b^m(\mathcal{T})$.) Then set

$$\overline{S}^m := \bigcup_{\mathcal{T} \in \mathbb{T}(\Gamma)} S^m(\mathcal{T}),$$

i.e., $\overline{S}^m$ is the collection of strategies consistent with m-level reasoning in some type structure. Write $\overline{S}_c^m := \text{proj}_{S_c} \overline{S}^m$ and observe that the set $\overline{S}_c^0, \ldots, \overline{S}_c^m, \ldots$ are decreasing. Then, for each finite m, $L_c^m = \overline{S}_c^m \backslash \overline{S}_c^{m+1}$, and $L_c^\infty = S_c \backslash \bigcup_{m \geq 0} L_c^m$. Moreover, $\overline{S}_c^\infty \subseteq L_c^\infty$.

In practice, there is a challenge in implementing this approach: We defined the set $\overline{S}_c^m$ as the union of the sets $S_c^m(\mathcal{T})$ across all type structures $\mathcal{T}$ associated with the given game Γ. Observe that this step requires searching across all type structures $\mathcal{T} \in \mathbb{T}(\Gamma)$. But, for a given finite extensive-form game, $\mathbb{T}(\Gamma)$ is uncountable. Therefore, this step involves an infinite task.

To overcome the issue, we will seek to identify $\overline{S}^m$ from properties of the game Γ alone. Theorem 6.1 will provide a method for doing so, provided the game is generic. Moreover, Proposition 6.1 will simplify the method provided by Theorem 6.1. (See also Section 8.1.)

[6] From this, we obtain $\mathcal{L}_c^*$: Put $d \in L_c^{*,m}$ if and only if $m = \max\{n : s_c \in L_c^n$ and there exists s_{-c} with $\delta(s_c, s_{-c}) = d\}$.

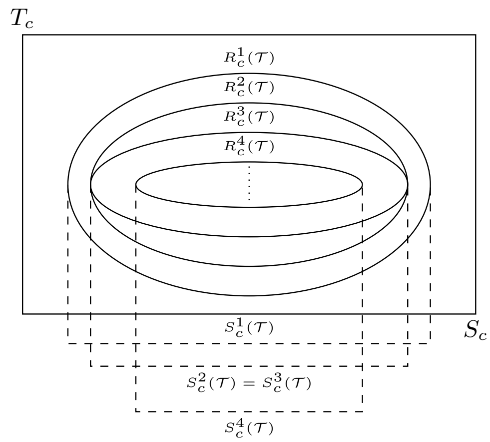

Figure 4.1: Projections of $R^m(\mathcal{T})$

5 Extensive-Form Best Response Property

In light of the above, we seek to identify properties based on the game Γ (alone) so that $Q_a \times Q_b \subseteq S_a \times S_b$ satisfies the properties if and only if, there exists a type structure $\mathcal{T} \in \mathbb{T}(\Gamma)$, so that

$$Q_a \times Q_b = S^m(\mathcal{T}) := \text{proj}_{S_a \times S_b} R^m(\mathcal{T}).$$

For the specific case of $m = \infty$ (i.e., RCSBR) such properties are known. Specifically, $Q_a \times Q_b$ is consistent with RCSBR if and only if it is an extensive-form best response set. (See Battigalli and Friedenberg, 2012.) With this in mind, we will use the solution concept of extensive-form best response sets as a benchmark, by which we think of addressing the broader issue.

5.1 Characterizing RCSBR

Given an array $\mathrm{p}_c \in \mathcal{A}(S_{-c}, \mathcal{E}_c)$, write $\mathbb{BR}[\mathrm{p}_c]$ for the set of strategies s_c that are a sequential best response under $\mathrm{p}_c \in \mathcal{A}(S_{-c}, \mathcal{E}_c)$.

Definition 5.1. Call $Q_a \times Q_b \subseteq S_a \times S_b$ an **extensive-form best response set** (**EFBRS**) if, for each $s_c \in Q_c$, there exists some CPS $\mathrm{p}_c \in \mathcal{C}(S_{-c}, \mathcal{E}_c)$ so that the following hold:

(1) $s_c \in \mathbb{BR}[\mathrm{p}_c]$,

(2) p_c strongly believes Q_{-c}, and

(3) if $r_c \in \mathbb{BR}[\mathrm{p}_c]$, then $r_c \in Q_c$.

In BoS with an Outside Option, there are three EFBRS's, viz. $\{Out\} \times \{L, R\}$, $\{Out\} \times \{R\}$, and $\{In\text{-}U\} \times \{L\}$.

Proposition 5.1 (Battigalli and Friedenberg, 2012)**.** *Fix a game* Γ*.*

(i) For each type structure $\mathcal{T}$*,* $\text{proj}_S R^\infty(\mathcal{T})$ *is an EFBRS.*

(ii) Given an EFBRS $Q_a \times Q_b$*, there exist a type structure* $\mathcal{T}$ *so that* $\text{proj}_S R^\infty(\mathcal{T}) = Q_a \times Q_b$*.*

This result says that EFBRS's characterize the behavior consistent with RCSBR. The EFBRS concept involves a fixed point requirement: If $Q_a \times Q_b$ is an EFBRS then, for each $s_a \in Q_a$, there is a CPS p_a under which s_a is optimal, so that p_a strongly believes the 'prediction' Q_b. This fixed-point requirement naturally arises from characterizing RCSBR behavior. To see this, fix some type structure $\mathcal{T}$ and a strategy-type pair (s_a, t_a) at which there is RCSBR. First, s_a is optimal under $\text{marg}_{S_b}\beta_a(t_a)$. Second, since $\text{marg}_{S_b}\beta_a(t_a)$ assigns probability one to each $\text{proj}_{S_b} R_b^m(\mathcal{T})$, $\text{marg}_{S_b}\beta_a(t_a)$ assigns probability one to $\text{proj}_{S_b} R_b^\infty(\mathcal{T})$. Third, if r_a is also optimal under $\text{marg}_{S_b}\beta_a(t_a)$, then (r_a, t_a) is also rational. It then follows that $(r_a, t_a) \in R_a^m(\mathcal{T})$ for all m, i.e., $r_a \in \text{proj}_{S_a} R_a^\infty(\mathcal{T})$.

Proposition 5.1 allows us to identify the sets $S_a^\infty(\mathcal{T}) \times S_b^\infty(\mathcal{T})$ based on properties of the game alone. As a consequence:

Corollary 5.1. *For each game* Γ*,*

$$\overline{S}^\infty := \bigcup_{\mathcal{T} \in \mathbb{T}(\Gamma)} S^\infty(\mathcal{T}) = \bigcup_{Q_a \times Q_b \text{ is an EFBRS}} (Q_a \times Q_b).$$

5.2 The m-BRP

A natural starting point is to turn this fixed-point definition into an iterative definition. Toward that end, fix a type structure $\mathcal{T}$ and observe that the sets $R^0(\mathcal{T}), R^1(\mathcal{T}), R^2(\mathcal{T}), \ldots$ satisfy the following:

(1) Each $\text{proj}_S R^m(\mathcal{T})$ is a product set, i.e., $\text{proj}_S R^m(\mathcal{T}) = \text{proj}_{S_a} R^m(\mathcal{T}) \times \text{proj}_{S_b} R^m(\mathcal{T})$.

(2) The sequence $(\text{proj}_S R^0(\mathcal{T}), \text{proj}_S R^1(\mathcal{T}), \ldots)$ is decreasing, i.e., $\text{proj}_S R^{m+1}(\mathcal{T}) \subseteq \text{proj}_S R^m(\mathcal{T})$.

With this in mind, we will focus on a sequence of strategy profiles that satisfy a decreasing property. Say $(Q^0, \ldots, Q^m)$ is a **decreasing sequence of strategy profiles** if (i) $Q^0 = S_a \times S_b$, (ii) each $Q^n = Q_a^n \times Q_b^n$ is a product set, and (iii) for each $n = 0, \ldots, m-1$, $Q^{n+1} \subseteq Q^n$.

Definition 5.2. Say $X = X_a \times X_b$ satisfies the **(extensive-form) best response property relative to** $(\mathbf{Q^0}, \ldots, \mathbf{Q^m})$ if $(Q^0, \ldots, Q^m, X)$ is a decreasing sequence of strategy profiles satisfying the following property: For each $s_c \in X_c$, there exists a CPS $\mathrm{p}_c \in \mathcal{C}(S_{-c}, \mathcal{E}_c)$ so that

(BRP.1) $s_c \in \mathbb{BR}[\mathrm{p}_c]$,

(BRP.2) p_c strongly believes $Q_{-c}^0, \ldots, Q_{-c}^m$, and

(BRP.3) if $r_c \in \mathbb{BR}[\mathrm{p}_c]$, then $r_c \in X_c$.

Definition 5.3. Let $m \geq 1$. Say $(Q^0, \ldots, Q^m)$ satisfies the **m-(extensive-form) best response property (m-BRP)** if

(i) Q^1 is non-empty and

(ii) for each $n = 1, \ldots, m-1$, Q^{n+1} satisfies the best response property relative to $(Q^0, \ldots, Q^n)$.

The m-BRP is a natural analogue of the EFBRS, converting that definition from a fixed-point definition to an iterative definition: If $s_c \in Q_c^n$, then s_c is sequentially optimal under an array that strongly believes lower-order prediction sets $Q_{-c}^0, \ldots, Q_{-c}^{n-1}$, instead of the prediction of the same order Q_{-c}^n. This is reflected by conditions (BRP.1)-(BRP.2). Condition (BRP.3) is the same maximality property of the EFBRS. (See Example 5.1 in Battigalli and Friedenberg (2009) for an explanation of the maximality criterion.)

Remark 5.1. Fix $m \geq 2$ and $(Q^0, \ldots, Q^m)$. Then $(Q^0, \ldots, Q^m)$ satisfies the m-BRP if and only if (i) $(Q^0, \ldots, Q^{m-1})$ satisfies the $(m-1)$-BRP and (ii) Q^m satisfies the extensive-form best response property relative to $(Q^0, \ldots, Q^{m-1})$. □

5.3 Characterizing RmSBR: Challenges

Say Q is **consistent with the m-BRP** if there exists some $(m-1)$-BRP, viz. $(Q^0, \ldots, Q^{m-1})$, so that Q satisfies the extensive-form best response property relative to $(Q^0, \ldots, Q^{m-1})$. The natural conjecture is that we can characterize the set $\overline{S}^m$ as the union over the sets Q that are consistent with the m-BRP, i.e.,

$$\overline{S}^m = \bigcup_{Q \text{ is consistent with the m-BRP}} Q. \tag{1}$$

If so, we would have a method to compute the set $\overline{S}^m$, without making reference to type structures.

This section is devoted to exploring the extent to which the conjecture is vs. is not true. First, we begin with a benchmark result (the proof of which is standard). The benchmark result is a step toward showing the conjecture. We then point to challenges for establishing the conjecture. The next section provides the characterization for a generic class of games.

We begin with a benchmark result: The RmSBR predictions induce a sequence of sets that satisfies the $(m+1)$-BRP.

Proposition 5.2. *Fix a type structure $\mathcal{T}$. The sequence $(\text{proj}_S R^0(\mathcal{T}), \ldots, \text{proj}_S R^m(\mathcal{T}))$ satisfies the m-BRP.*

Proposition 5.2 implies that each $S^m(\mathcal{T}) \equiv \text{proj}_S R^m(\mathcal{T})$ is consistent with the m-BRP. As such, the left-hand side of Equation (1) is contained in the right-hand side of Equation (1). The question is whether the converse obtains: For each set Q consistent with the m-BRP, does there exists some $\mathcal{T}$ so that $Q \subseteq \text{proj}_S R^m(\mathcal{T})$?

One might conjecture that, for each m-BRP $(Q^0, \ldots, Q^m)$, there exists a type structure $\mathcal{T}$ so that

$$Q^n \subseteq \text{proj}_S R^n(\mathcal{T}) \qquad \text{for all } n = 1, \ldots, m. \tag{2}$$

If so, it would suffice to deliver the result. However, as the next series of examples show, this conjecture is incorrect.

Example 5.1. Consider the game in Figure 5.1. Let (Q^0, Q^1, Q^2) be the decreasing sequence of strategy profiles with

$$Q_a^1 \times Q_b^1 = S_a \times \{y_1q_1, y_1q_2, y_2\} \qquad \text{and} \qquad Q_a^2 \times Q_b^2 = \{x_2\} \times Q_b^1$$

Note this is a 2-BRP.[7] But, we show that there is no type structure $\mathcal{T}$ with $Q^1 \subseteq \text{proj}_S R^1(\mathcal{T})$ and $Q^2 = \text{proj}_S R^2(\mathcal{T})$. Thus, Equation (2) cannot be strengthened to equality.

Suppose otherwise. Then there exists a type t_a so that $(x_1z_1, t_a) \in R_a^1(\mathcal{T})$. Observe that $\beta_a(t_a)$ must assign probability one to $\{y_2\} \times T_b$ at each information set. But, y_2 is a sequential best response under

[7] Let us point to several features of the example: First, x_1z_1 and x_2 are both optimal under a CPS that assigns probability 1 to y_2, x_1z_2 is uniquely optimal under a CPS that assigns probability 1 to y_3, and x_2 is uniquely optimal under a CPS that assigns probability 1 to $\{y_1q_1, y_1q_2\}$. Second, y_1q_1 (resp. y_1q_2) and y_2 are the *only* strategies that are optimal under a CPS that assigns probability 1 to x_2 at the initial information set and then assigns probability one to x_1z_2 (resp. x_1z_1) conditional upon observing x_1. Third, y_2 is uniquely optimal under a CPS that assigns probability 1 to x_1z_2.

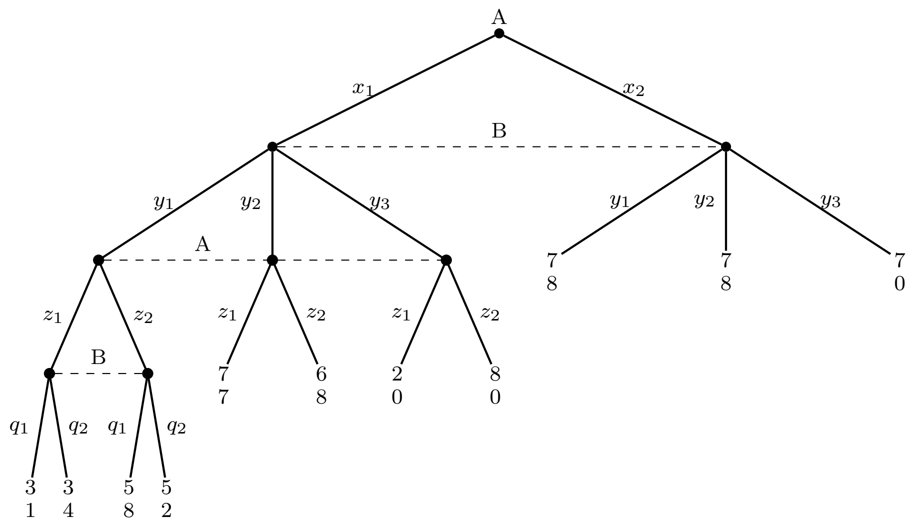

Figure 5.1

each CPS and, so, $\{y_2\} \times T_b \subseteq R_b^1(\mathcal{T})$. With this, t_a strongly believes $R_b^1(\mathcal{T})$ and so $(x_1 z_1, t_a) \in R_a^2(\mathcal{T})$. Thus, $Q_a^2 \neq R_a^2(\mathcal{T})$. □

Example 5.1 shows that we may have a 2-BRP (Q^0, Q^1, Q^2) so that, there is no type structure $\mathcal{T}$, with both $Q^1 = \text{proj}_S R^1(\mathcal{T})$ and $Q^2 = \text{proj}_S R^2(\mathcal{T})$. But, this is immaterial from the perspective of delivering the desired result—i.e., from the perspective of delivering Equation (1). This is because there is a type structure $\mathcal{T}$, with both $Q^1 \subseteq \text{proj}_S R^1(\mathcal{T})$ and $Q^2 \subseteq \text{proj}_S R^2(\mathcal{T})$. In fact, this conclusion holds more generally:

Proposition 5.3. *Fix a game* Γ.

(i) For each 1-BRP (Q^0, Q^1)*, there exists some* $\mathcal{T}$ *so that* $Q^1 = \text{proj}_S R^1(\mathcal{T})$.

(ii) For each 2-BRP (Q^0, Q^1, Q^2)*, there exists some* $\mathcal{T}$ *so that* $Q^1 = \text{proj}_S R^1(\mathcal{T})$ *and* $Q^2 \subseteq \text{proj}_S R^2(\mathcal{T})$.

In light of Proposition 5.3, Equation (1) does indeed hold for $m = 1, 2$. However, we will next see that an analogue of Proposition 5.3 does not hold for the 3-BRP.

Example 5.2. Return to the game in Figure 5.1. Let (Q^0, Q^1, Q^2, Q^3), where (Q^0, Q^1, Q^2) is the 2-BRP described in Example 5.1 and

$$Q_a^3 \times Q_b^3 = Q_a^2 \times \{y_1 q_1, y_2\}.$$

We will show that there is no type structure $\mathcal{T}$ so that $Q^n \subseteq \text{proj}_S R^n(\mathcal{T})$ for each $n = 1, 2, 3$.

Suppose, contra hypothesis, that such a type structure $\mathcal{T}$ exists. Since $Q^3 \subseteq \text{proj}_S R^3(\mathcal{T})$, there exists some t_b with $(y_1 q_1, t_b) \in R_b^3(\mathcal{T})$. Then, $\beta_b(t_b)$ must assign positive probability to $\{x_1 z_2\} \times T_a$ conditional upon $\{x_1 z_1, x_1 z_2\} \times T_a$. We will argue that $(\{x_1 z_2\} \times T_a) \cap R_a^2(\mathcal{T}) = \emptyset$ but $(\{x_1 z_1\} \times T_a) \cap R_a^2(\mathcal{T}) \neq \emptyset$, contradicting the fact that t_b strongly believes $R_a^2(\mathcal{T})$.

First, observe that $(x_1 z_1) \in Q_a^1$ and so, by assumption, $(x_1 z_1) \in \text{proj}_{S_a} R_a^1(\mathcal{T})$. Thus, repeating the argument in Example 5.1 above, $(x_1 z_1) \in \text{proj}_{S_a} R_a^2(\mathcal{T})$. Second, observe that $x_1 z_2$ is only a sequential best response under a CPS that assigns positive probability to $\{y_3\} \times T_b$ at the initial information set. Since y_3 is dominated, no such CPS can strongly believe $R^1_b(\mathcal{T})$. Thus, $x_1 z_2 \notin \text{proj}_{S_a} R^2_a(\mathcal{T})$. □

Let us review Example 5.2. It gives a 3-BRP so that, if $Q^1 \subseteq \text{proj}_S R^1(\mathcal{T})$, then there exists some strategy in Q^3 that is not contained in $\text{proj}_S R^3(\mathcal{T})$. This key is that there is a strategy in Q^3_b that is sequentially optimal under a CPS that strongly believes Q^2_a. But, that CPS cannot strongly believe $\text{proj}_{S_a} R^2_a(\mathcal{T})$; this arises because Q^2_a must be a strict subset of $\text{proj}_{S_a} R^2_a(\mathcal{T})$.

With this in mind, the next section restricts attention to a class of games that are generic. In that class, we can guarantee that a strong variant of Equation (5.1): If $(Q^0, Q^1, \ldots, Q^m)$ is an m-BRP, then we can find a type structure $\mathcal{T}$ so that $Q^n = \text{proj}_S R^n(\mathcal{T})$ for each $n = 1, \ldots, m$.

6 Characterizing RmSBR: Generic Games

Say two strategies s_c and r_c are **equivalent** if they induce the same plan of action, i.e., $\zeta(s_c, \cdot) = \zeta(r_c, \cdot)$. Write $[s_c]$ for the set of strategies that are equivalent to s_c, and observe that, since the game is non-trivial, each $[s_c] \subsetneq S_c$. So, if s_c and r_c are equivalent, then $\pi_c(s_c, \cdot) = \pi_c(r_c, \cdot)$. It follows that $s_c \in \mathbb{BR}[\mathrm{p}_c]$ if and only if $[s_c] \subseteq \mathbb{BR}[\mathrm{p}_c]$.[8]

Definition 6.1. Call a game Γ **generic** if the following property holds: If there exists a CPS $\mathrm{p}_c \in \mathcal{C}(S_{-c}, \mathcal{E}_c)$ so that $s_c \in \mathbb{BR}[\mathrm{p}_c]$, then there exists a CPS $\mathrm{q}_c \in \mathcal{C}(S_{-c}, \mathcal{E}_c)$ so that $[s_c] = \mathbb{BR}[\mathrm{q}_c]$.

Thus, a game is generic if any sequentially optimal strategy is "uniquely" sequentially optimal under some (perhaps different) CPS. Here, unique is taken to mean "up to equivalent strategies." Note, Example 5.1 is non-generic. The strategy $x_1 z_1$ is only sequentially optimal under a CPS that assigns probability 1 to y_2 at the initial information set. However, x_2 is *also* sequentially optimal under that CPS. In 6.2 below, we will point to classes of games that satisfy this genericity requirement. We first point out that, in generic games, RmSBR is characterized by the $(m+1)$-BRP.

6.1 Characterization

When a game is generic, in the sense of this paper, the predictions of RmSBR are exactly captured by the sets consistent with the $(m+1)$-BRP:

Theorem 6.1. *Fix a generic game Γ. The following hold for each m.*

(i) For each epistemic game $\mathcal{T}$, $(\text{proj}_S R^0(\mathcal{T}), \ldots, \text{proj}_S R^m(\mathcal{T}))$ satisfies the m-BRP.

(ii) If $(Q^0, \ldots, Q^m)$ satisfies the m-BRP, then there exists some type structure $\mathcal{T}$ so that $(\text{proj}_S R^0(\mathcal{T}), \ldots, \text{proj}_S R^m(\mathcal{T})) = (Q^0, \ldots, Q^m)$.

Part (i) is a special case of Proposition 5.2. Part (ii) is specific to generic games. It says that, for a generic game and an associated m-BRP, we can construct a type structure so that, for each $n = 0, \ldots, m-1$, the predictions of RnSBR are exactly captured by Q^{n+1}.

Let us provide a sketch of the proof of Theorem 6.1(ii). The aim is to highlight the role of genericity. Throughout, fix a generic game Γ and a 2-BRP (Q^0, Q^1, Q^2). The goal is to construct a type structure $\mathcal{T}$

[8] In BoS with an outside option, *O-L* and *O-R* are two equivalent strategies. We have simply been writing *Out*; our notation formally reflects an equivalence class of strategies.

so that $\text{proj}_S R^1 = Q^1$ and $\text{proj}_S R^2 = Q^2$. (Note, here and in the sketch below, we surpress reference to $\mathcal{T}$. No confusion should result.)

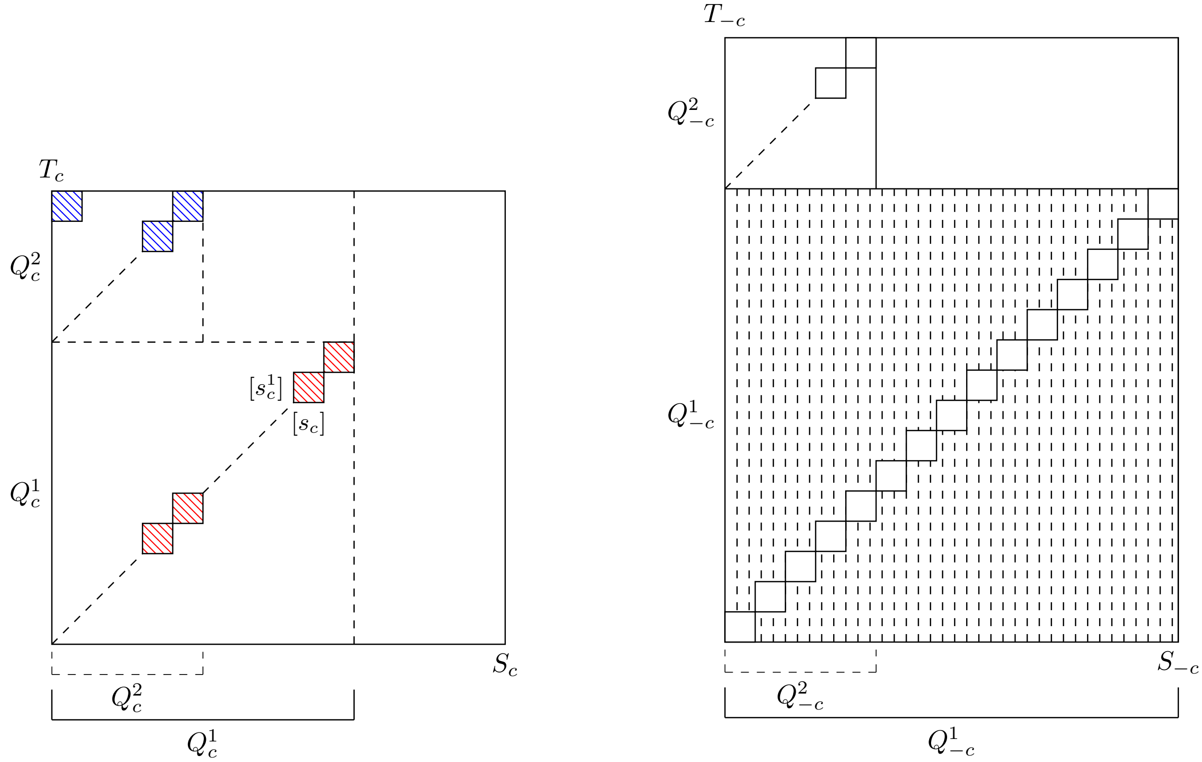

(a) Type Sets

(b) Belief $\beta_c(t_c)(\cdot|S_{-c}(h) \times T_{-c})$ when $Q_{-c}^1 = S_{-c}$

Figure 6.1: Construction of Type Structure

Begin by setting the type set for player c to be $T_c = Q_c^1 \bigsqcup Q_c^2$. For each $s_c \in Q_c^2 \subseteq Q_c^1$, there will be two associated types: a 1-type labeled s_c^1 and a 2-type labeled s_c^2. The idea will be to construct belief maps so that 1-strategy-type pairs (s_c, s_c^1) are rational but do not strongly believe rationality, and 2-strategy-type pairs (s_c, s_c^2) are rational and strongly believe rationality. In fact, we will ask for somewhat more. Refer to Figure 6.1a. We will seek to construct belief maps satisfying two requirements: First, $R_c^1 \backslash R_c^2$ is 'along the diagonal' of $Q_c^1 \times Q_c^1$ modulo equivalence classes, i.e.,

$$R_c^1 \backslash R_c^2 = \{[s_c] \times [s_c^1] : s_c \in Q_c^1\}.$$

This diagonal is illustrated as the red squares in Figure 6.1a. Second, R_c^2 contains the diagonal of $Q_c^2 \times Q_c^2$ and is contained in the square $Q_c^2 \times Q_c^2$. The blue rectangles in Figure 6.1a illustrate this set.

Begin by constructing the beliefs associated with 2-types. By definition of a 2-BRP, for each $s_c \in Q_c^2$, there is a CPS $j_c(s_c^2)$ on $(S_{-c}, \mathcal{E}_c)$ so that $[s_c] \subseteq \mathbb{BR}[j_c(s_c^2)] \subseteq Q_c^2$ and $j_c(s_c^2)$ strongly believes Q_{-c}^1. Choose $\beta_c(s_c^2)$ so that (i) $\text{marg}_{S_{-c}} \beta_c(s_c^2) = j_c(s_c^2)$ and (ii) $\beta_c(s_c^2)((s_{-c}, s_{-c}^1)|S_{-c}(h) \times T_{-c}) = j_c(s_c^2)(s_{-c}|S_{-c}(h))$, whenever $s_{-c} \in S_{-c}(h) \cap Q_{-c}^1$. So, if $S_{-c}(h) \cap Q_{-c}^1 \neq \emptyset$, then $\beta_c(s_c^2)(\cdot|S_{-c}(h) \times T_{-c})$ is concentrated on the diagonal of $Q_{-c}^1 \times Q_{-c}^1$.

Next construct the beliefs associated with 1-types. For the purpose of illustrating the construction, we focus on two extreme cases: where Q^1_{-c} is a singleton (i.e., $Q^1_{-c} = [s^*_{-c}]$) and where Q^1_{-c} is the full strategy set (i.e., $Q^1_{-c} = S_c$). In both cases, for each $s_c \in Q^1_c$, there is a CPS $j_c(s^1_c)$ on $(S_{-c}, \mathcal{E}_c)$ so that $[s_c] = \mathbb{BR}[j_c(s^1_c)]$. In the first case, we can and do choose this CPS so that $j_c(s^1_c)$ does not strongly believe Q^1_{-c}. (Lemma C.3 in the Appendix shows how this is done.) We then choose $\beta_c(s^1_c)$ so that $\text{marg}_{S_{-c}} \beta_c(s^1_c) = j_c(s^1_c)$ and each $\beta_c(s^1_c)(\cdot|S_{-c}(h) \times T_{-c})$ is concentrated on the diagonal of $Q^1_{-c} \times Q^1_{-c}$, when possible. In the second case, we instead adopt an anti-diagonal construction. Refer to Figure 6.1b. Observe that each $\beta_c(s^2_c)(\cdot|S_{-c}(h) \times T_{-c})$ is concentrated off the diagonal of $Q^1_{-c} \times Q^1_{-c}$. Specifically, if $s_{-c} \in [t_{-c}] \subseteq Q^1_{-c}$, then $\beta_c(s^2_c)((s_{-c}, t_{-c})|S_{-c}(h) \times T_{-c}) = 0$.

Observe that, under the construction, we have

$$R^1_c = \{[s_c] \times [s^1_c] : s_c \in Q^1_c\} \cup \bigcup_{s^2_c \in Q^2_c} (\mathbb{BR}[j_c(s^2_c)] \times \{s^2_c\}).$$

Now, by the diagonal construction, types in Q^2_c strongly believe R^1_{-c}. Moreover, types in Q^1_c do not strongly believe R^1_{-c}. (If Q^1_{-c} is a singleton, this follows from the fact that $j_c(s^1_c)$ does strongly believe Q^1_{-c}. If $Q^1_{-c} = S_{-c}$ this follows from the anti-diagonal construction.) Thus, $R^2_c = \bigcup_{s^2_c \in Q^2_c} (\mathbb{BR}[j_c(s^2_c)] \times \{s^2_c\})$.

6.2 Generic Games

Genericity is a property stated in terms of sequential best responses. We now identify classes of games that satisfy genericity. A natural starting point is games that satisfy no relevant ties (Battigalli, 1997):

Definition 6.2. A game satisfies **no relevant ties** if $\pi_c(s_c, s_{-c}) = \pi_c(r_c, s_{-c})$ implies $\zeta(s_c, s_{-c}) = \zeta(r_c, s_{-c})$.

More informally, a game satisfies no relevant ties if, whenever player c is decisive over two distinct terminal nodes z and z^* (i.e., if there exists (s_c, s_{-c}) and (r_c, s_{-c}) with $\zeta(s_c, s_{-c}) \neq \zeta(r_c, s_{-c})$), she is not indifferent between those terminal nodes.

A perfect-information game satisfying no relevant ties is generic. (See Lemma D.8.) However, a non-perfect information game may satisfy no relevant ties, even though it is not generic. (Appendix D provides an example.) That is, beyond perfect-information games, no relevant ties is not sufficient for the game to be generic. With this in mind, we will introduce a property called no relevant convexities. (Appendix D notes that the condition can be weakened). This property implies no relevant ties, but is sufficient for any finite game of perfect recall to be generic. Moreover, as we will see, if the game satisfies no relevant convexities, then we will be able to simplify the procedure for identification.

Definition 6.3. Fix some $X = X_c \times X_{-c} \subseteq S_c \times S_{-c}$ and some $s_c \in X_c$. Say r_c **supports** s_c **with respect to** $X_c \times X_{-c}$ if there exists $\sigma_c \in \mathcal{P}(S_c)$ with $\text{Supp}\, \sigma_c \subseteq X_c$ such that

(1) $r_c \in \text{Supp}\, \sigma_c$, and

(2) $\pi_c(\sigma_c, s_{-c}) = \pi_c(s_c, s_{-c})$ for all $s_{-c} \in X_{-c}$.

A strategy r_c supports s_c on $X_c \times X_{-c}$, if it is in the support of a convex combination $\sigma_c \in \mathcal{P}(X_c)$, so that the convex combination is payoff equivalent to s_c on X_{-c}. Thus, in a sense, if r_c supports s_c on $X_c \times X_{-c}$, then r_c can be viewed as "indifferent" to s_c on X_{-c}. No relevant convexities requires that, in that case, s_c and r_c reach the same terminal node, whenever a strategy $s_{-c} \in X_{-c}$ is played.

Definition 6.4. The game satisfies **no relevant convexities** (**NRC**) if the following holds: If r_c supports s_c with respect to some $S_c \times X_{-c} \subseteq S_c \times S_{-c}$, then $\zeta(s_c, s_{-c}) = \zeta(r_c, s_{-c})$ for each $s_{-c} \in X_{-c}$.

The formal results only require a weaker version of NRC, described in Appendix D. Note that if a game satisfies no relevant convexities, then it also satisfies no relevant ties: If it were to fail no relevant ties, there would be some $\zeta(s_c, s_{-c}) \neq \zeta(r_c, s_{-c})$ with $\pi_c(s_c, s_{-c}) = \pi_c(r_c, s_{-c})$. It follows that $s_c, r_c \in S_c$, r_c supports s_c with respect to some $S_c \times \{s_{-c}\}$, but $\zeta(s_c, s_{-c}) \neq \zeta(r_c, s_{-c})$. So, it also fails no relevant convexities.

If a game satisfies no relevant convexities, then it is generic. (See Corollary D.2.) Thus, Theorem 6.1 will apply. In fact, when the game satisfies NRC, we can simplify the definition of the m-BRP, in a way that will serve useful for computational purposes.

Proposition 6.1. *Fix a game that satisfies NRC. Then $(Q^0, \ldots, Q^m)$ satisfies the m-BRP if and only if Q^1 is non-empty and, for each $n = 1, \ldots, m$ and each $s_c \in Q^n_c$, there exists an array $\mathrm{p}_c \in \mathcal{A}(S_{-c}, \mathcal{E}_c)$ with*

(i) $s_c \in \mathbb{BR}[\mathrm{p}_c]$,

(ii) p_c *strongly believes* $Q^0_{-c}, \ldots, Q^{n-1}_{-c}$, *and*

(iii) $[s_c] \subseteq Q^n_c$.

Proposition 6.1 says that, when the game satisfies NRC, we can simplify the m-BRP definition in two ways. First, we can replace the maximality criterion (BRP.3) with the requirement that does not depend on the CPS's: We now simply require that, if $s_c \in Q^m_c$, then every strategy equivalent to s_c is also in Q^m_c. Second (and, arguably, more importantly) we can replace CPSs with arrays. This will prove useful, from the perspective of computation. Appendix D shows that, when the game fails NRC, we may not be able to replace CPS's with arrays.

7 Illustration of Identification: Centipede Games

To illustrate the approach to identification, we apply the analysis to the Centipede game. Refer to Figure 7.1: We order the non-terminal nodes (or vertices) as $1, 2, \ldots, V$, where $V \geq 3$. (So, 1 indicates the initial node and V indicates the last non-terminal vertex.) The payoffs are as specified in the figure with $x, y > 0$. The figure depicts the game when the number of non-terminal nodes (or vertices) V is odd; when the number of non-terminal nodes (or vertices) V is instead even, Bob moves at V. In that case, the payoffs from out_V are $(x + (V-3)y, x + Vy)$ and the payoffs from in_V are $(x + Vy, x + (V-1)y)$.[9]

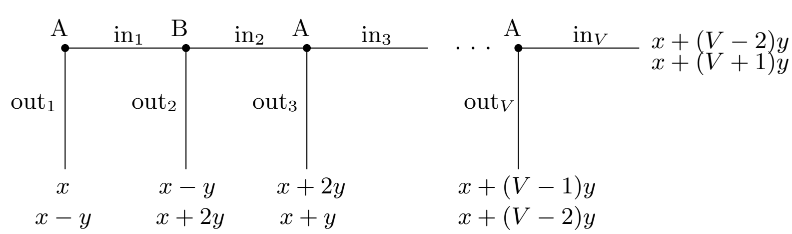

Figure 7.1: Centipede Game

Let us point to three crucial features of the game: First, the player who moves at vertex $v \leq V - 1$ strictly prefers out_{v+2} (resp. in_V if $v = V - 1$) to out_v and strictly prefers out_v to out_{v+1}. Second, the

[9]The first component in the payoff vector reflects Ann's payoffs.

player who moves at vertex V, strictly prefers out_V to in_V. Finally, the game has no relevant ties and, so, is generic. Thus, we will be able to exploit Theorem 6.1 to identify reasoning in the game.

Write $[\text{out}, v]_c$ for the set of strategies of player c which play $\text{in}_{v'}$ at vertices that v' (strictly) precede v and out_v. Write $[\text{in}]_c$ for the set that contains the (unique) strategy of of player c, which specifies $\text{in}_{v'}$ at every node v'.

A Benchmark A useful benchmark will be m rounds of EFR. As in Example 3.2, write EFR_c^m for the m-EFR strategies for player c. The m-EFR strategy set will depend on whether m is odd versus even and whether player c moves last (i.e., whether V is odd versus even). With this in mind, we take the notational convention that, if V is odd then player ℓ (player 'last') is Ann and if V is even then player ℓ is Bob. We write $-\ell$ for the second-to-last player.

With this, EFR is characterized as follows:

- $\text{EFR}_\ell^1 \times \text{EFR}_{-\ell}^1 = (S_\ell \backslash [\text{in}]_\ell) \times S_{-\ell}$,
- $\text{EFR}_\ell^2 \times \text{EFR}_{-\ell}^2 = \text{EFR}_\ell^1 \times (S_{-\ell} \backslash [\text{in}]_{-\ell})$,
- $\text{EFR}_\ell^m \times \text{EFR}_{-\ell}^m = (\text{EFR}_\ell^{m-1} \backslash [\text{out}, V+3-m]_\ell) \times \text{EFR}_{-\ell}^{m-1}$ if $m = 3, \ldots, V$ is odd, and
- $\text{EFR}_\ell^m \times \text{EFR}_{-\ell}^m = \text{EFR}_\ell^{m-1} \times (\text{EFR}_{-\ell}^{m-1} \backslash [\text{out}, V+3-m]_{-\ell})$ if $m = 4, \ldots, V$ is even.

For all $m \geq V+1$, $\text{EFR}_\ell^m \times \text{EFR}_{-\ell}^m = \text{EFR}_\ell^4 \times \text{EFR}_{-\ell}^4$. Note, this also corresponds round-for-round with the backward induction algorithm.

The m-BRP Procedure The m-BRP procedure has very different implications for the first-mover (Ann) and the second-mover (Bob). In particular, for the second-mover, it allows for 'many predictions.' Specifically:

Remark 7.1.

(i) If V is odd, then $\overline{S}_b^m = S_b$ for all $m \geq 1$.

(ii) If V is even, then for $\overline{S}_b^m = (S_b \backslash [\text{in}]_b)$ for all $m \geq 1$.

To understand this remark, note that, when V is odd (resp. even) $[\text{out}, 1]_a \times S_b$ (resp. $[\text{out}, 1]_a \times (S_b \backslash [\text{in}]_b)$) is an EFBRS. Thus, for each m, it is contained in the union over the m-BRPs.

The situation is quite different for the first-player:

Proposition 7.1. *In the centipede game, for each m, $\overline{S}_a^m = \text{EFR}_a^m$.*

At first glance, this result might appear trivial: For each m, $\text{EFR}_a^m \times \text{EFR}_b^m$ is consistent with the m-BRP. Thus, $\text{EFR}_a^m \subseteq \overline{S}_a^m$. However, the key is showing that $\overline{S}_a^m \subseteq \text{EFR}_a^m$ and, as we have seen, this is not the case for the second-mover b. (Appendix E explains why it is true.)

Implications for Identifying Reasoning in Games The implications for identifying reasoning are quite stark. Without imposing auxiliary assumptions about players' beliefs, the researcher cannot identify the second player's reasoning. But, the researcher can identify the first-player's reasoning.

To better understand these points, consider a lab experiment, in which the researcher has subjects choose strategies via the strategy method. (This will allow the researcher to observe the full strategy of player i, even if the other player's behavior prevents player i's information set from being reached.) The researcher can have subjects play centipede games of varying lengths. By observing the behavior only in the role of Ann (i.e., the first player), the researcher can identify the maximum level of reasoning consistent with observed behavior (for any given subject). Specifically, for any m, there is some centipede game that will allow the researcher to infer whether the subject's behavior in the role of Ann is consistent with R(m-1)SBR. By observing behavior across centipede games of varying length, the researcher can separate R(m-1)SBR from RmSBR.[10]

That said, absent auxiliary assumptions about beliefs, the researcher cannot use the data in the role of Bob (i.e., the second player) to identify reasoning: If V is odd, then all observed behavior by Bob is consistent with rationality and common strong belief of rationality. If V is even, then playing $[\text{in}]_b$ would identify irrationality; playing any other strategy would again be consistent with rationality and common strong belief of rationality. Thus, identifying Bob's behavior as a Level-1 Reasoner requires imposing auxiliary assumptions about Bob's belief conditional upon an irrational move by Ann.

Remark 7.2. A word of caution: One might conclude that a computer can/should take the role of Bob. However, to the extent that the subjects' reasoning about rationality is not entirely determined by limits on ability (Friedenberg, Kets and Kneeland, 2016), doing so may change *how* a subject reasons about rationality. That is, it may very well be important to have real subjects play in the role of the second player, even if the researcher will disregard their play. □

8 An Algorithm for Identification

Refer to Theorem 6.1 and Section 4.2: In a generic games, we can identify the maximum m consistent with RmSBR provided we can identify all sets consistent with the $(m+1)$-BRP. This provides a procedure for identifying reasoning in dynamic games.

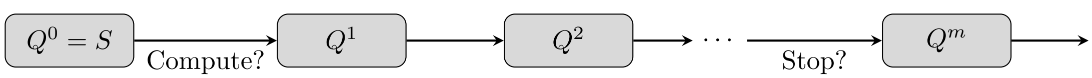

Figure 8.1: m-BRP Elimination Procedure

From a computational perspective, there are two obstacles that hinder implementing the procedure. Refer to Figure 8.1: First, to determine if (Q^0, Q^1) is an 1-BRP, we need to compute the set of all $\mathbb{BR}_c[\mathrm{p}_c]$ for all CPS's p_c. However, there are potentially uncountably many such CPS's. And, analogously, for any higher-order m-BRPs. Second, we must determine that the procedure stops. Because the game is finite, there must exist some M so that $Q^M = Q^m$ for all $m \geq M$, i.e., where the m-BRP stops shrinking. However, from the perspective of implementing the identification procedure, the researcher must know *when* it stops shrinking. We will see that this step is not obvious.

This section addresses both computational issues. In so doing, it provides an algorithm for identification—at least for a wide class of games.

[10] A potential concern would be learning—i.e., if the subject learns how to play the game between different centipede games. However, this concern can be mitigated by adopting some of the experimental design features in Kneeland (2015). Specifically, the experimenter can give no feedback and allow the subject to 'change play' after reviewing the choices across a series of centipede games.

8.1 Computing Best Responses

Fix some decreasing sequence of strategies $(Q^0, \ldots, Q^m)$. We seek to identify the sets $Q = Q_a \times Q_b$ that satisfy the best response property relative to $(Q^0, \ldots, Q^m)$. Toward that end, we will restrict attention to games that satisfy NRC. In that case, we can identify all the sets Q by a simple test on the strategies in Q^m.

Specifically, a strategy $s_c \in Q_c^m$ will "pass the test" if there is an array $\mathrm{p}_c \in \mathcal{A}(S_{-c}, \mathcal{E}_c)$ so that $s_c \in \mathbb{BR}_c[\mathrm{p}_c]$ and p_c strongly believes $Q^0_{-c}, \ldots, Q^m_{-c}$. When the game satisfies NRC, a set Q satisfies the best response property relative to $(Q^0, \ldots, Q^m)$ if each Q_c can be written as a union over sets $[s_c] \subseteq Q_c^m$ so that s_c passes the test.

Fix some $h \in H_c$ and write $n(h) = \max\{n : Q^n_{-c} \cap S_{-c}(h) \neq \emptyset\}$. Then, enumerate

$$Q^{n(h)}_{-c} \cap S_{-c}(h) = \{s^1_{-c}, \ldots, s^K_{-c}\} \qquad \text{and} \qquad S_c(h) = \{s^1_c, \ldots, s^L_c\}.$$

Say a strategy $s_c \in Q_c^m$ **passes the test at** h if either $s_c \notin S_c(h)$ or, there exists $(\mu_1, \ldots, \mu_K) \geq (0, \ldots, 0)$ with $\sum_{k=1}^K \mu^k = 1$, so that s_c maximizes $\sum_{k=1}^K \pi_c(\cdot, s^k_{-c})\mu^k$ amongst all strategies in $S_c(h) = \{s^1_c, \ldots, s^L_c\}$. A strategy s_c **passes the test** if it passes the test at each $h \in H_c$.

The key is that the simplex algorithm can be used to determine if s_c passes the test at h. Specifically, when $s_c \in S_c(h)$, the problem is equivalent to choosing $(\mu^1, \ldots, \mu^K, \sigma^1, \ldots, \sigma^L)$ to solve

$$\begin{aligned}
\text{Maximize} \quad & \sum_{k=1}^K \pi_c(s_c, s^k_{-c})\mu^k \\
& \text{subject to} \quad \textstyle\sum_{k=1}^K [\pi_c(s_c, s^k_{-c}) - \pi_c(s^l_c, s^k_{-c})]\mu^k + \sigma^l = 0 \quad \text{for each } l = 1, \ldots, L \\
& \qquad\qquad\quad\; \mu^1 + \mu^2 + \cdots + \mu^K = 1 \\
& \qquad\qquad\quad\; (\mu^1, \ldots, \mu^K, \sigma^1, \ldots, \sigma^L) \geq (0, \ldots, 0)
\end{aligned}$$

We can apply the simplex algorithm to this linear programming problem. The algorithm terminates by either (a) concluding that there is no feasible solution, (b) providing an optimal solution, or (c) concluding that the objective function is unbounded over the feasible region. (See Chapter 2 in Bradley, Hax and Magnanti, 1977.) In the first scenario s_c fails the test and in the latter two scenarios s_c passes the test.

8.2 Termination of the Procedure

Fix some decreasing sequence of strategies $(Q^0, Q^1, Q^2, \ldots)$, where each $(Q^0, \ldots, Q^m)$ satisfies the m-BRP. Since each $Q^{m+1} \subseteq Q^m$, we can think of $(Q^0, Q^1, Q^2, \ldots)$ as defining an elimination procedure. Since the strategy set is finite, this procedure must terminate—i.e., there exists some M so that, for each $m \geq M$, $Q^m = Q^M$. If the researcher knew at which M this occurred, the researcher could use that fact to determine that the elimination has stopped.

At first glance, there might appear to be straightforward route to determine this M. Typically, an elimination procedure stops shrinking at the first round where no strategy is eliminated for either player. However, this same principle does not apply to the m-BRP elimination procedure. We may have $Q^m = Q^{m-1}$ even though $Q^{m+1} \subsetneq Q^m$.

Example 8.1. Consider the simultaneous-move game given by Figure 8.2 and note that the game is generic. Yet, for each m, there is an m-BRP with $(Q^0, \ldots, Q^m)$, so that (i) for each $n \leq m$, $Q^n = \{U, D\} \times \{L, R\}$,

		Bob	
		L	R
Ann	U	1,1	0,0
	D	0,0	1,1

Figure 8.2

and (ii) $Q^{m+1} = \{U, D\} \times \{R\}$. Thus, the $(m+1)$-BRP procedure has no shrinkage up until round m, but a shrinkage at round $(m+1)$. $\square$

To understand why Example 8.1 can occur, refer to Figure 4.1. We can have $\text{proj}_S R^3 = \text{proj}_S R^2$, even though $R^3(\mathcal{T}) \subsetneq R^2(\mathcal{T})$. Example 8.1 highlights the fact that we can have arbitrarily long pauses before shrinkage: For any M, we can construct some m-BRP $(Q^0, \ldots, Q^M, \ldots, Q^m)$ so that the $(Q^0, Q^1, Q^2, \ldots)$ procedure has not terminated within M steps.

Nonetheless, we can provide a bound on the elimination procedure $(\overline{S}^0, \overline{S}^1, \overline{S}^2, \ldots)$, i.e., we can find some M so that, for all $m \geq M$, $\overline{S}^m = \overline{S}^M$. To understand why, consider an m-BRP procedure $(Q^0, Q^1, Q^2, \ldots)$ with a pause at round m, i.e., $Q^{m+1} = Q^m$ but $Q^n \subsetneq Q^{m+1}$ for some $n \geq m+2$. The key is that any eliminated strategy—i.e., any strategy in $Q_c^{m+1} \backslash Q_c^n$—must be contained in $\overline{S}^n$. That is, there must exist some other m-BRP procedure $(\hat{Q}^0, \hat{Q}^1, \hat{Q}^2, \ldots)$ so that $Q^{m+1} \backslash Q^n \subseteq \hat{Q}^n$. This follows from the following observation.

Observation 8.1. *Fix some $(Q^0, Q^1, Q^2, \ldots)$ where, for each m, $(Q^0, \ldots, Q^m)$ is an m-BRP. If $Q^{m+1} = Q^m$, then Q^m is an EFBRS.*

Thus, if $Q^{m+1} = Q^m$, then we can define $(\hat{Q}^0, \hat{Q}^1, \hat{Q}^2, \ldots)$ with, for each $n \geq 1$, $\hat{Q}^n = Q^m$. This defines a new m-BRP procedure. From this, it follows that $Q^m = \hat{Q}^n \subseteq \overline{S}^n$ for all $n \geq 1$.

Proposition 8.1. *Fix a game Γ and set*

$$\overline{M} = \begin{cases} 2\min\{|S_a|, |S_b|\} - 1 & \text{if } |S_a| \neq |S_b|, \\ 2\min\{|S_a|, |S_b|\} - 2 & \text{if } |S_a| = |S_b|. \end{cases}$$

Then, for all $m \geq \overline{M}$, $\overline{S}^m = \overline{S}^\infty$.

Proposition 8.1 provides a bound $\overline{M}$ for the procedure $(\overline{S}^0, \overline{S}^1, \overline{S}^2, \ldots)$. Thus, it suffices to compute all the $\overline{M}$-BRPs, $(Q^0, \ldots, Q^{\overline{M}})$.

In practice, it is often not necessary to compute all the $\overline{M}$-BRPs. Refer to Figure 8.3. We begin with $Q^0 = S$ and identify all the 1-BRPs (Q^0, Q^1). We use these 1-BRP's to identify all the 2-BRPs (Q^0, Q^1, Q^2). And so on. We know we can stop after we have identified all the $\overline{M}$-BRPs. However, along any given "BRP path" $(Q^0, Q^1, \ldots, Q^{\overline{M}})$, it may be possible to stop prior to round $\overline{M}$. Specifically, it may be possible to stop at $m < \overline{M}$ if $Q^m = Q^{m+1}$.

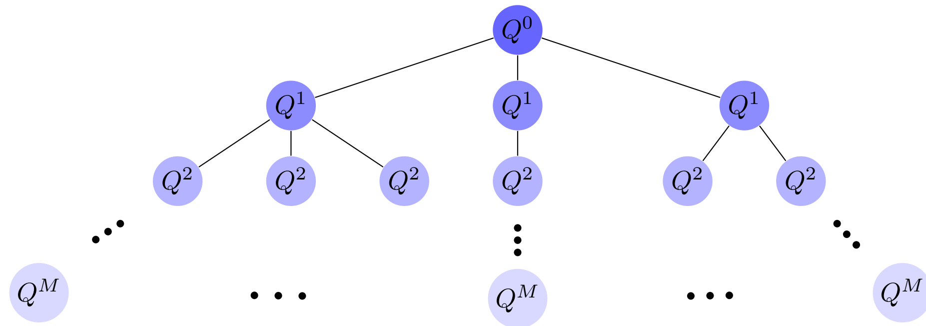

Figure 8.3: m-BRP Elimination Tree

9 Discussion

a. Bounded Reasoning about Rationality This paper focuses on bounded reasoning about rationality. With this in mind, we define Ann to be a Level-m Reasoner if her behavior is consistent with R(m-1)SBR but inconsistent with RmSBR. Notice, we take no position on *why* Ann might exhibit such bounded reasoning about rationality. For instance, Ann may face such bounds if she has interacted in the past with a population of like-minded individuals and has observed behavior that leads her to bet on "irrational" behavior. Or, alternatively, Ann may face such bounds if she faces limitations in her ability to engage in interactive reasoning—i.e., if she finds it difficult to specify what she thinks "Bob thinks she thinks etc..." about behavior.

At first glance, this latter interpretation may seem inconsistent with our analysis: Type structures induce infinite hierarchies of beliefs about the play of the game—i.e., m^{th}-order beliefs about play, for all m. However, this contradiction is illusory. The key observation is that, if Ann is a Level-m Reasoner (as defined in this paper), then hierarchies of beliefs beyond level m do not affect her behavior. Formally, consider two types t_a and u_a with the same $(m+1)^{th}$-order beliefs about the strategies played in the game. For any strategy s_a, the strategy-type pair (s_a, t_a) is consistent with RmSBR if and only if (s_a, u_a) is consistent with RmSBR. The higher-order beliefs become an artifact of our formalism and do not have any behavioral significance.[11]

b. Strong Belief versus Initial Belief Example 3.1 explained why we focus on "strong belief of rationality" and not "full belief of rationality." An alternate is to focus instead on "initial belief of rationality" (Ben-Porath, 1997). A type initially believes rationality if, at the start of the game, the type assigns probability one to the event that the other player is rational.

Initial belief relaxes what it means to reason about rationality throughout the game. In so doing, it allows us to rationalize the data at higher levels: If a state is consistent with rationality and m^{th}-order strong belief of rationality, then the state is consistent with rationality and m^{th}-order initial belief of rationality. Thus, if a subject is identified as being a Level-m Reasoner according to "strong belief" and a Level-n Reasoner according to "initial belief," then $n \geq m$.

The implication is that, under initial belief, it is more difficult to identify levels of reasoning. For instance, refer back to the three-legged Centipede game. There, both $[\text{out}]_1$ and $[\text{out}]_3$ are consistent with

[11] So, in particular, we could instead adapt the type structure frameworks in Kets (2010) and Heifetz and Kets (N.d.) (which allow for finite-order beliefs) to conditional beliefs; in doing so, we would reach analogous conclusions.

"rationality and common initial belief of rationality." Thus, under initial belief, we could not use the first player's behavior to identify levels of reasoning, absent imposing auxiliary assumptions.

c. Common Reasoning Remark 4.1 pointed out a subtlety in the definition of a Level-∞ Reasoner: A strategy s_c may be consistent with RmSBR for all m (i.e., for each m, there may be some type structure $\mathcal{T}^m$ and associated $(s_c, t_c^m) \in R_c^m(\mathcal{T}^m)$). But, this does not immediately imply that s_c is consistent with RCBR (i.e., there exists some type structure $\mathcal{T}$ so that $(s_c, t_c) \in \bigcap_m R_c^m(\mathcal{T})$ for some type t_c). However, in light of Observation 8.1, we can conclude that Level-∞ reasoning and common reasoning are in fact equivalent.

Corollary 9.1. *The following are equivalent:*

(i) For each m, there exists some $\mathcal{T}^m$ so that $s_c \in \text{proj}_{S_c} R_c^m(\mathcal{T}^m)$.

(ii) There exists some $\mathcal{T}$ so that $s_c \in \text{proj}_{S_c} \bigcap_m R_c^m(\mathcal{T})$.

Proof. It is immediate that (ii) implies (i). Suppose that (i) hold and choose $M \geq 2\min\{|S_a|, |S_b|\}$. Then there exists some $m \leq M-1$ so that

$$\text{proj}_S R^m(\mathcal{T}^M) = \text{proj}_S R^{m+1}(\mathcal{T}^M).$$

It follows from Proposition 5.2 and Observation 8.1 that $\text{proj}_S R^m(\mathcal{T}^M)$ is an EFBRS. Thus, by Proposition 5.1, there exists some $\mathcal{T}$ so that

$$\text{proj}_S \bigcap_{m \geq 1} R^m(\mathcal{T}) = \text{proj}_S R^m(\mathcal{T}^M).$$

It follows that

$$\text{proj}_{S_c} R_c^M(\mathcal{T}^M) \subseteq \text{proj}_S R^m(\mathcal{T}^M) = \text{proj}_S \bigcap_{m \geq 1} R^m(\mathcal{T}),$$

as desired. ■

With this, the data is consistent with Level-∞ Reasoning for c if and only if the data is consistent with common reasoning for c.

d. Simultaneous Move Games Example 5.1 showed that we may have a 2-BRP (Q^0, Q^1, Q^2) so that there is no type structure $\mathcal{T}$ with $Q^1 = \text{proj}_S R^1(\mathcal{T})$ and $Q^2 = \text{proj}_S R^2(\mathcal{T})$. This can also be the case for a simultaneous move game.[12] However, if (Q^0, Q^1, Q^2) is a 2-BRP, there exists a type structure $\mathcal{T}$ with $Q^1 = \text{proj}_S R^1(\mathcal{T})$ and $Q^2 \subseteq \text{proj}_S R^2(\mathcal{T})$. (See Proposition 5.3.) In simultaneous move games, this generalizes beyond the 2-BRP. Specifically, if $(Q^0, \ldots, Q^m)$ is a m-BRP, there exists a type structure $\mathcal{T}$ with each $Q^n \subseteq \text{proj}_S R^n(\mathcal{T})$. As such, in any simultaneous move game, $\overline{S}^m$ is the union over all Q consistent with the m-BRP and this set coincides with the set of m-rationalizable strategies. (This fits with well-known results.)

With this in mind, observe that there is no analogue of Example 5.2 for simultaneous move games. The reason is because, in simultaneous move games, strong belief is monotonic, whereas, in dynamic games,

[12] An example is available upon request.

strong belief is non-monotonic.[13] Thus, in a simultaneous move game, if a CPS strongly believes Q^2_{-c} and $Q^2_{-c} \subsetneq \text{proj}_{S_{-c}} R^2_{-c}(\mathcal{T})$, then the CPS also strongly believes $\text{proj}_{S_{-c}} R^2_{-c}(\mathcal{T})$. However, if in a dynamic game, a CPS may strongly believe some Q^2_{-c} with $Q^2_{-c} \subsetneq \text{proj}_{S_{-c}} R^2_{-c}(\mathcal{T})$, even though the CPS does not strongly believe $\text{proj}_{S_{-c}} R^2_{-c}(\mathcal{T})$.

e. Beyond Generic Games Example 5.2 illustrates that, for a given m-BRP $(Q^0, Q^1, \ldots, Q^m)$, there may not be a type structure $\mathcal{T}$ so that $Q^n \subseteq \text{proj}_S R^n(\mathcal{T})$ for each $n = 1, \ldots, m$. This can arise because the game is non-generic. The example leaves open that there may be an alternate m-BRP $(\hat{Q}^0, \hat{Q}^1, \ldots, \hat{Q}^m)$ so that the following hold:

(i) $\hat{Q}^m = Q^m$, and

(ii) there exists some type structure $\mathcal{T}$ so that $\hat{Q}^n \subseteq \text{proj}_S R^n(\mathcal{T})$ for each $n = 1, \ldots, m$.

If correct, it would say that Equation (2) holds for all games. We neither know this to be true nor have a counterexample. Thus, we leave it as an open question.

f. Two versus Three Player Games We have restricted attention to two player games. When the game has three (or more) players, two conceptual questions arise. First, do players have independent or correlated beliefs about their co-players? (See Brandenburger and Friedenberg, 2008 on this issue.) Second, do players engage in correlated versus independent rationalization? (See Section 9c in Battigalli and Friedenberg, 2012 on this issue.) Our analysis applies to the n-player game verbatim, provided players have correlated beliefs and engage in correlated rationalization.

Appendix A Preliminaries

This appendix provides preliminary results, which are used in subsequent results.

Marginalization Property of Belief

Lemma A.1. *Fix epistemic game $\mathcal{T}$. If $\beta_c(t_c)$ strongly believes the event $E_{-c} \subseteq S_{-c} \times T_{-c}$, then* $\text{marg}_{S_{-c}} \beta_c(t_c)$ *strongly believes* $\text{proj}_{S_{-c}} E_{-c}$.

Proof. Suppose $\beta_c(t_c)$ strongly believes the event $E_{-c} \subseteq S_{-c} \times T_{-c}$. Fix some $S_{-c}(h) \times T_{-c} \in \mathcal{E}_c \otimes T_{-c}$. If $\text{proj}_{S_{-c}} E_{-c} \cap S_{-c}(h) \neq \emptyset$, then there exists $(s_{-c}, t_{-c}) \in E_{-c}$ so that $s_{-c} \in S_{-c}(h)$. It follows that $E_{-c} \cap (S_{-c}(h) \times T_{-c}) \neq \emptyset$ and so $\beta_c(E_{-c}|S_{-c}(h) \times T_{-c}) = 1$. Now note that

$$\text{marg}_{S_{-c}} \beta_c(\text{proj}_{S_{-c}} E_{-c}|S_{-c}(h) \times T_{-c}) = \beta_c(\text{proj}_{S_{-c}} E_{-c} \times T_{-c}|S_{-c}(h) \times T_{-c}) \geq \beta_c(E_{-c}|S_{-c}(h) \times T_{-c}).$$

It follows that $\text{marg}_{S_{-c}} \beta_c(\text{proj}_{S_{-c}} E_{-c}|S_{-c}(h) \times T_{-c}) = 1$, as desired. ■

Image CPS's: Fix a CPS $p_c \in \mathcal{C}(S_{-c}, \mathcal{E}_c)$ and some (measurable) mapping $\tau_{-c} : S_{-c} \to S_{-c} \times T_{-c}$. Define q_c as follows: For each conditioning event $S_{-c}(h) \times T_{-c} \in \mathcal{E}_c \otimes T_{-c}$, set

$$q_c(E_{-c}|S_{-c}(h) \times T_{-c}) = p_c((\tau_{-c})^{-1}(E_{-c})|S_{-c}(h))$$

[13] In simultaneous move games, strong belief coincides with "belief," i.e., *ex ante* assigning probability one to an event.

for each Borel $E_{-c} \subseteq S_{-c} \times T_{-c}$. We refer to q_c at the **image CPS of** p_c **under** τ_{-c}. So defined, q_c is indeed a CPS. See Battigalli, Friedenberg and Siniscalchi (2012, Part III, Chapter 4). Moreover, if $\tau_{-c}(s_{-c}) \in \{s_{-c}\} \times T_{-c}$ for each s_{-c}, then the image CPS of p_c under τ_{-c}, viz. q_c, has $\mathrm{marg}_{S_{-c}} q_c = p_c$. As a consequence, for any given CPS $\mathrm{p}_c \in \mathcal{C}(S_{-c}, \mathcal{E}_c)$, we can find some CPS $q_c \in \mathcal{C}(S_{-c} \times T_{-c}, \mathcal{E}_c \otimes T_{-c})$ so that $\mathrm{marg}_{S_{-c}} q_c = p_c$.

Structure of Games and Sequential Best Responses By perfect recall, we have the following: (i) For each $h, h' \in H_c$, either $S(h) \subseteq S(h')$, $S(h') \subseteq S(h)$, or $S(h) \cap S(h') = \emptyset$. (ii) For each $h \in H_c$, $S(h) = S_c(h) \times S_{-c}(h)$. The second of these implies the following:

Lemma A.2. *Fix $h, h' \in H_c$ so that $S(h) \cap S(h') = \emptyset$. If $S_{-c}(h) \cap S_{-c}(h') \neq \emptyset$, then $S_c(h) \cap S_c(h') = \emptyset$.*

Proof. Fix $h, h' \in H_c$ so that $S_c(h) \cap S_c(h') \neq \emptyset$ and $S_{-c}(h) \cap S_{-c}(h') \neq \emptyset$. Then there exists some $s_c \in S_c(h) \cap S_c(h')$ and $s_{-c} \in S_{-c}(h) \cap S_{-c}(h')$. It follows that $(s_c, s_{-c}) \in S_c(h) \times S_{-c}(h)$ and $(s_c, s_{-c}) \in S_c(h') \times S_{-c}(h')$. By perfect recall, $S(h) = S_c(h) \times S_{-c}(h)$ and $S(h') = S_c(h') \times S_{-c}(h')$. Thus, $S(h) \cap S(h') \neq \emptyset$. ■

Lemma A.3. *Fix $h^*, h^{**} \in H_c$ so that $S(h^{**}) \subseteq S(h^*)$. Let $p_c \in \mathcal{P}(S_{-c})$ with $p_c(S_{-c}(h^*)) = 1$ and $p_c(S_{-c}(h^{**})) > 0$. If $s_c \in S_c(h^{**})$ is optimal under p_c given all strategies in $S_c(h^*)$, then s_c is optimal under $p_c(\cdot|S_{-c}(h^{**}))$ given all strategies in $S_c(h^{**})$.*

Proof. Suppose that there exists some $r_c \in S_c(h^{**})$ so that

$$\sum_{s_{-c}} [\pi_c(r_c, s_{-c}) - \pi_c(s_c, s_{-c})] p_c(s_{-c}|S_{-c}(h^{**})) > 0.$$

Construct a strategy $\tilde{r}_c$ so that

$$\tilde{r}_c(h) = \begin{cases} r_c(h) & \text{if } S(h) \subseteq S(h^{**}) \\ s_c(h) & \text{otherwise.} \end{cases}$$

Fix some $s_{-c} \in S_{-c}(h^{**})$ and observe that (s_c, s_{-c}) and (r_c, s_{-c}) are both contained in $S(h^{**}) = S_c(h^{**}) \times S_{-c}(h^{**})$. (This follows from perfect recall.) Thus, $(\tilde{r}_c, s_{-c}) \in S(h^{**})$ and so $\tilde{r}_c \in S_c(h^{**}) \subseteq S_c(h^*)$.

We will show that

(i) $\zeta(r_c, s_{-c}) = \zeta(\tilde{r}_c, s_{-c})$ if $s_{-c} \in S_{-c}(h^{**})$, and

(ii) $\zeta(s_c, s_{-c}) = \zeta(\tilde{r}_c, s_{-c})$ if $s_{-c} \in S_{-c}(h^*) \backslash S_{-c}(h^{**})$.

From this, it follows that

$$\sum_{s_{-c}} [\pi_c(\tilde{r}_c, s_{-c}) - \pi_c(s_c, s_{-c})] p_c(s_{-c}) > 0.$$

contradicting the hypothesis that s_c is optimal under p_c given all strategies in $S_c(h^*)$.

First, fix some $s_{-c} \in S_{-c}(h^{**})$ and note that, by perfect recall,

$$(s_c, s_{-c}), (r_c, s_{-c}), (\tilde{r}_c, s_{-c}) \in S_c(h^{**}) \times S_{-c}(h^{**}) = S(h^{**}).$$

Suppose, contra hypothesis, that $\zeta(r_c, s_{-c}) \neq \zeta(\tilde{r}_c, s_{-c})$. Then there exists some $h \in H_c$ so that $(r_c, s_{-c}), (\tilde{r}_c, s_{-c}) \in S(h) = S_c(h) \times S_{-c}(h)$ but $r_c(h) \neq \tilde{r}_c(h) = s_c(h)$. By construction, $\neg(S(h) \subseteq S(h^{**}))$. Since $S(h^{**}) \cap$

$S(h) \neq \emptyset$, it follows that $S(h^{**}) \subsetneq S(h)$. Thus we have established that $s_c(h) \neq r_c(h)$ and $(s_c, s_{-c}), (r_c, s_{-c}) \in S(h^{**})$; but this contradicts perfect recall.

Second, fix some $s_{-c} \in S_{-c}(h^*) \backslash S_{-c}(h^{**})$ and suppose, contra hypothesis, that $\zeta(s_c, s_{-c}) \neq \zeta(\tilde{r}_c, s_{-c})$. Then there exists some $h \in H_c$ with $(s_c, s_{-c}), (\tilde{r}_c, s_{-c}) \in S(h) = S_c(h) \times S_{-c}(h)$ and $s_c(h) \neq \tilde{r}_c(h) = r_c(h)$. By construction, $S(h) \subseteq S(h^{**})$, contradicting the assumption that $s_{-c} \in S_{-c}(h^*) \backslash S_{-c}(h^{**})$. ■

Appendix B Appendix for Section 5.3

We begin with the proof of Proposition 5.2.

Proof of Proposition 5.2. We will show that, for each $m \geq 1$, $(\text{proj}_S R^0(\mathcal{T}), \ldots, \text{proj}_S R^m(\mathcal{T}))$ satisfies the m-BRP. The proof is by induction on m.

$\mathbf{m = 1}$: If $s_c \in \text{proj}_{S_c} R^1_c$, then there exists some $t_c \in T_c$ so that $(s_c, t_c) \in R^1_c(\mathcal{T})$. Take $p_c = \text{marg}_{S_{-c}} \beta_c(t_c)$. Note that $s_c \in \mathbb{BR}[\text{p}_c]$. Moreover, if $r_c \in \mathbb{BR}[\text{p}_c]$, then $(r_c, t_c) \in R^1_c$ and so $r_c \in \text{proj}_{S_c} R^1_c$.

$\mathbf{m \geq 2}$: Assume the claim holds for m and fix some $(\text{proj}_S R^0(\mathcal{T}), \ldots, \text{proj}_S R^m(\mathcal{T}), \text{proj}_S R^{m+1}(\mathcal{T}))$. Then, by the induction hypothesis, $(\text{proj}_S R^0(\mathcal{T}), \ldots, \text{proj}_S R^m(\mathcal{T}))$ satisfies the m-BRP. Thus, it suffices to show that $\text{proj}_S R^{m+1} = \text{proj}_{S_a} R^{m+1} \times \text{proj}_{S_b} R^{m+1}$ satisfies the extensive-form best response property relative to $(\text{proj}_S R^0(\mathcal{T}), \ldots, \text{proj}_S R^m(\mathcal{T}))$.

Fix some $s_c \in \text{proj}_{S_c} R^{m+1}(\mathcal{T})$. There exists some $t_c \in T_c$ so that $(s_c, t_c) \in R^{m+1}_c(\mathcal{T})$. Take $p_c = \text{marg}_{S_{-c}} \beta_c(t_c)$. Since $(s_c, t_c) \in R^1_c(\mathcal{T})$, $s_c \in \mathbb{BR}[\text{p}_c]$. Moreover, $\beta_c(t_c)$ strongly believes $R^0_{-c}(\mathcal{T}), \ldots, R^m_{-c}(\mathcal{T})$. So applying Lemma A.1, $\text{marg}_{S_{-c}} \beta_c(t_c)$ strongly believes $\text{proj}_{S_{-c}} R^0_{-c}(\mathcal{T}), \ldots, \text{proj}_{S_{-c}} R^m_{-c}(\mathcal{T})$. Finally, if $r_c \in \mathbb{BR}[\text{p}_c]$, then $(r_c, t_c) \in R^{m+1}_c(\mathcal{T})$ and so $r_c \in \text{proj}_{S_c} R^{m+1}_c(\mathcal{T})$. ■

We next prove Proposition 5.3. We divide the proofs into two parts.

Proof of Proposition 5.3(i). Fix a 1-BRP (Q^0, Q^1). Construct $\mathcal{T}$ as follows: Set $T_c = Q^1_c$. For each $s_c \in T_c = Q^1_c$, choose $\beta_c(s_c)$ so that $\text{marg}_{S_c} \beta_c(s_c)$ is a CPS p_c with $[s_c] \in \mathbb{BR}[\text{p}_c] \subseteq Q^1_c$. (The fact that such a CPS exists follows from the definition of a 1-BRP.) It follows that $\text{proj}_{S_c} R^1_c(\mathcal{T}) = Q^1_c$. ■

Proof of Proposition 5.3(ii). Fix a 2-BRP (Q^0, Q^1, Q^2). For each $s_c \in Q^1_c$, there exists some CPS $j_c[s_c]$ so that $s_c \in \mathbb{BR}[j_c[s_c]] \subseteq Q^1_c$. Moreover, if $s_c \in Q^2_c$, we can take $j_c[s_c]$ to strongly believe Q^1_{-c} and so that $\mathbb{BR}[j_c[s_c]] \subseteq Q^2_c$.

With this in mind, set $T_c = Q^1_c$ and define $\beta_c(s_c)$ as follows: For each $s_{-c} \in S_{-c}(h) \cap Q^1_{-c}$,

$$\beta_c(s_c)((s_{-c}, s_{-c})|S_{-c}(h) \times T_{-i}) = j_c[s_c](s_{-c}|S_{-c}(h)).$$

For each $s_{-c} \in S_{-c}(h) \cap (S_{-c} \backslash Q^1_{-c})$,

$$\beta_c(s_c)(\{s_{-c}\} \times T_{-c}|S_{-c}(h) \times T_{-i}) = j_c[s_c](s_{-c}|S_{-c}(h)).$$

Then,

$$R^1_c(\mathcal{T}) = \bigcup_{s_c \in Q^1_c} (\mathbb{BR}[j_c[s_c]] \times \{s_c\})$$

and so $\text{proj}_{S_c} R^1_c(\mathcal{T}) = Q^1_c$. Moreover, if $s_c \in Q^2_c$, type s_c strongly believes $R^1_{-c}(\mathcal{T})$. As such, $Q^2_c \subseteq \text{proj}_{S_c} R^2_c(\mathcal{T})$. ■

Appendix C Proof of Theorem 6.1

To show Theorem 6.1, it will be useful to introduce a strong justification property. With this in mind, refer to a set $X_c \subseteq Q_c$ as an **effective singleton** if there exists some s_c so that $X_c = [s_c]$. If $X_c \subseteq Q_c$ is not effectively a singleton, then we simply say it is **non-singleton**.

Definition C.1. Fix an m-BRP $(Q^0, \ldots, Q^m)$. Say that the m-BRP satisfies the **strong justification property** if, for each player c and each $n = 1, \ldots, m$, we can find a mappings $j^n_c : Q^n_c \to \mathcal{C}(S_{-c}, \mathcal{E}_c)$ satisfying the following criteria:

(j.a) For each $s_c \in Q^1_c$, $\mathbb{BR}[j^1_c(s_c)] = [s_c]$. Moreover, if Q^1_{-c} if effectively a singleton, then $j^1_c(s_c)$ does not strongly believe Q^1_{-c}.

(j.b) For each $n = 2, \ldots, m$ and each $s_c \in Q^n_c$, $s_c \in \mathbb{BR}[j^n_c(s_c)] \subseteq Q^n_c$ and $j^n_c(s_c)$ strongly believes $Q^0_{-c}, \ldots, Q^{n-1}_{-c}$.

Observe that, by definition of an m-BRP, we can always find mappings $j^n_c : Q^n_c \to \mathcal{C}(Q_{-c}, \mathcal{E}_c)$ satisfying condition (j.b). But, condition (j.a) is stronger than that required by an m-BRP. If we find mappings $j_c = (j^1_c, \ldots, j^m_c)$ satisfying these requirements, we say that j_c **strongly justifies the m-BRP for player** c or j_a **and** j_b **strongly justify the m-BRP**.

Theorem 6.1 follows from the following two propositions.

Proposition C.1. *Fix an m-BRP $(Q^0, \ldots, Q^m)$ satisfying the strong justification property. Then there exists an associated epistemic game $\mathcal{T}$ so that, for each $n = 1, \ldots, m$,* $\text{proj}_S R^n(\mathcal{T}) = Q^n$.

Proposition C.2. *If the game is generic, then any m-BRP satisfies the strong justification property.*

We now turn to proving these two results.

C.1 Proof of Proposition C.1

Throughout we fix an m-BRP $(Q^0, \ldots, Q^m)$ satisfying the strong justification property. Thus, for each player c, there are mappings $j_c = (j^1_c, \ldots, j^m_c)$ that strongly justify the m-BRP.

Description of the Type Structure For each player c and each $n = 1, \ldots, m$, set $U^m_c \equiv Q^m_c$ and write $v^n_c : Q^n_c \to U^n_c$ for the identity map. The type set for player c will be $T_c = \bigsqcup^m_{n=1} U^n_c$. We will refer to types in U^n_c as the n-types for player c.

It will be convenient to specify the **diagonal** of $Q^n_c \times U^n_c$. This will be given by

$$\text{diag}^n_c = \bigcup_{s_c \in Q^n_c} ([s_c] \times v^n_c([s_c]))$$

Observe that, if $[s_c] = [r_c]$ then $\upsilon_c^n([s_c]) = \upsilon_c^n([r_c])$ and so $[s_c] \times \upsilon_c^n([r_c]) \subseteq \text{diag}_c^n$. Moreover, if Q_c^n is non-singleton then, for each $s_c \in Q_c^n$, there exists a type $t_c \in U_c^n$ so that $(s_c, t_c) \in (Q_c^n \times U_c^n) \backslash \text{diag}_c^n$.

For each $n = 1, \ldots, m$, we will define a mapping $\tau_{-c}^n : S_{-c} \to S_{-c} \times T_{-c}$ with $\tau_{-c}^n(s_{-c}) \in \{s_{-c}\} \times T_{-c}$. In addition, the mappings will satisfy the following properties: For $n = 1$, if Q_{-c}^1 is non-singleton, then the range of τ_{-c}^1 is concentrated on $S_{-c} \times U_{-c}^1$ but off of diag_{-c}^1, i.e., each $\tau_{-c}^1(s_{-c}) \in (S_{-c} \times U_{-c}^1) \backslash \text{diag}_{-c}^1$. For $n = 2, \ldots, m$, for each $s_{-c} \in Q_{-c}^1$, $\tau_{-c}^n(s_{-c})$ is in the maximal diagonal consistent with s_{-c}. Specifically, for a given $s_{-c} \in Q_{-c}^1$, let $\ell = \max\{k = 1, \ldots, n-1 : s_{-c} \in Q_{-c}^k\}$ and set $\tau_{-c}^n(s_{-c}) = (s_{-c}, \upsilon_{-c}^{\ell}(s_{-c}))$.

The belief map is such that, for each $\upsilon_c^n(s_c) \in U_c^n$, $\beta_c(\upsilon_c^n(s_c))$ is the image CPS of $j_c^n(s_c)$ under τ_{-c}^n. Observe that, for each $s_c \in Q_c^n$, $\text{marg}_{Q_{-c}} \beta_c(\upsilon_c^n(s_c)) = j_c^n(s_c)$.

Analysis It will be convenient to define sets of n-strategy-type pairs of the players. In particular, for each player c and each $n = 1, \ldots, m$, set

$$\mathbb{Q}_c^n = \bigcup_{s_c \in Q_c^n} (\mathbb{BR}[j_c^n(s_c)] \times \{\upsilon_c^n(s_c)\}).$$

By Conditions (j.a)-(j.b) of strong justification, $\text{diag}_c^n \subseteq \mathbb{Q}_c^n$.

Lemma C.1. *For each* $n = 1, \ldots, m$, $\text{proj}_{S_c} \mathbb{Q}_c^n = Q_c^n$.

Proof. If $s_c \in Q_c^n$, then $s_c \in \mathbb{BR}[j_c^n(s_c)]$ and so $(s_c, \upsilon_c^n(s_c)) \in \mathbb{Q}_c^n$. Fix some $(s_c, \upsilon_c^n(r_c)) \in \mathbb{Q}_c^n$. Then, $r_c \in Q_c^n$ and $s_c \in \mathbb{BR}[j_c^n(r_c)]$. It follows that $s_c \in \mathbb{BR}[j_c^n(r_c)] \subseteq Q_c^n$, as required. ■

Lemma C.2. *For each* $n = 1, \ldots, m$, $R_a^n(\mathcal{T}) \times R_b^n(\mathcal{T}) = \bigcup_{k=n}^m (\mathbb{Q}_a^k \times \mathbb{Q}_b^k)$.

Proof. The case of $n = 1$ is immediate from the construction. Thus, we show $n = 2, \ldots, m$. The proof is by induction on n. In the proof, we write R_c^n instead of $R_c^n(\mathcal{T})$ since the type structure $\mathcal{T}$ is as constructed above.

Fix some $n = 2, \ldots, m$ and some $k = n, \ldots, m$ and some $(r_c, \upsilon_c^k(s_c)) \in \mathbb{BR}[j_c^k(s_c)] \times \{\upsilon_c^k(s_c)\} \subseteq \mathbb{Q}_c^k$. Since the claim holds for $n = 1$ it suffices to show the following:

(i) If $k = n - 1$, then $\upsilon_c^k(s_c)$ does not strongly believe R_{-c}^n.

(ii) If $k = n, \ldots, m$, then $\upsilon_c^k(s_c)$ strongly believes R_{-c}^n.

$\mathbf{n = 2}$: Fix some $k = 1, \ldots, m$ and some $(r_c, \upsilon_c^k(s_c)) \in \mathbb{BR}[j_c^k(s_c)] \times \{\upsilon_c^k(s_c)\} \subseteq \mathbb{Q}_c^k$. We will show (i)-(ii) hold. To do so, we will make use of the following property: Having established this claim for $n = 1$, we have that $R_{-c}^1 = \bigcup_{k=1}^m \mathbb{Q}_{-c}^k$ and so, by Lemma C.1, $Q_{-c}^1 = \text{proj}_{S_{-c}} \bigcup_{k=1}^m \mathbb{Q}_{-c}^k = \text{proj}_{S_{-c}} R_{-c}^1(\mathcal{T})$.

First, suppose that $k = 1$ and Q_{-c}^1 is an effective singleton. By Condition (j.a) of strong justification, $j_c^1(s_c)$ does not strongly believe Q_{-c}^1, i.e., there exists some information set h with $Q_{-c}^1 \cap S_{-c}(h) \neq \emptyset$ and $j_c^1(s_c)(S_{-c} \backslash Q_{-c}^1 | S_{-c}(h)) > 0$. Since $Q_{-c}^1 = \text{proj}_{S_{-c}} R_{-c}^1$, $R_{-c}^1 \cap (S_{-c}(h) \times T_{-c}) \neq \emptyset$. Moreover, $\beta_c(\upsilon_c^1(s_c))((S_{-c} \backslash Q_{-c}^1) \times T_{-c} | S_{-c}(h) \times T_{-c}) > 0$ and, again using the fact that $Q_{-c}^1 = \text{proj}_{S_{-c}} R_{-c}^1$, $((S_{-c} \backslash Q_{-c}^1) \times T_{-c}) \cap R_{-c}^1 = \emptyset$. Thus, $\upsilon_c^1(s_c)$ does not strongly believe R_{-c}^1.

Next, suppose that $k = 1$ and Q_{-c}^1 is non-singleton. Observe that, in this case,

$$\beta_c(\upsilon_c^1(s_c))(S_{-c} \times U_{-c}^1) \backslash \text{diag}_{-c}^1 | S_{-c} \times T_{-c}) = 1.$$

By Condition (j.a) of strong justification, if $(s_{-c}, t_{-c}) \in (S_{-c} \times U^1_{-c}) \backslash \text{diag}^1_{-c}$, then $s_c \notin \mathbb{BR}[j^1_{-c}(t_c)]$ and so $(s_{-c}, t_{-c}) \notin R^1_{-c}$. Thus, $\upsilon^1_c(s_c)$ does not strongly believe R^1_{-c}.

Finally, suppose that $k = 2, \ldots, m$. Fix a conditioning event $S_{-c}(h) \times T_{-c}$ so that $R^1_{-c} \cap (S_{-c}(h) \times T_{-c}) \neq \emptyset$. Since $Q^1_{-c} = \text{proj}_{S_{-c}} \mathbb{Q}^1_{-c} = \text{proj}_{S_{-c}} R^1_{-c}$, it follows that $Q^1_{-c} \cap S_{-c}(h) \neq \emptyset$. So, using the fact that $j^k_c(s_c)$ strongly believes Q^1_{-c}, it follows that $j^k_c(s_c)(Q^1_{-c}|S_{-c}(h)) = 1$. Now observe that, by construction,

$$\beta_c(\upsilon^k_c(s_c))(\bigcup_{l=1}^{k-1} \text{diag}^l_{-c}|S_{-c}(h) \times T_{-c}) = j^k_c(s_c)(Q^1_{-c}|S_{-c}(h)) = 1.$$

Since $\bigcup_{l=1}^{k-1} \text{diag}^l_{-c} \subseteq \bigcup_{l=1}^{m} \mathbb{Q}^{k-1}_{-c}$ and $R^1_{-c} = \bigcup_{l=1}^{m} \mathbb{Q}^l_{-c}$ (this result shown for $n = 1$), it follows that $\beta_c(\upsilon^k_c(s_c))(R^1_{-c}|S_{-c}(h) \times T_{-c}) = 1$, as desired.

$\mathbf{n \geq 3}$: Fix some $n = 3, \ldots, m$ and some $k = n, \ldots, m$ and some $(r_c, \upsilon^k_c(s_c)) \in \mathbb{BR}[j^k_c(s_c)] \times \{\upsilon^k_c(s_c)\} \subseteq \mathbb{Q}^k_c$. We show (i)-(ii).

First, suppose that $k = n - 1$. Fix (s_{-c}, t_{-c}) with $\beta_c(\upsilon^k_c(s_c))((s_{-c}, t_{-c})|S_{-c} \times T_{-c}) > 0$ and note that, by construction, $t_{-c} = \upsilon^{k-1}_{-c}(s_c)$. By the induction hypothesis (part (i)), $\upsilon^{k-1}_{-c}(s_c)$ does not strongly believe R^{n-2}_c. Thus, $\upsilon^n_c(s_c)$ does not strongly believe R^{n-1}_{-c}.

Second, suppose that $k = n, \ldots, m$. Fix a conditioning event $S_{-c}(h) \times T_{-c}$ so that $R^{n-1}_{-c} \cap (S_{-c}(h) \times T_{-c}) \neq \emptyset$. By the induction hypothesis and Lemma C.1

$$\text{proj}_{S_{-c}} R^{n-1}_{-c} = \text{proj}_{S_{-c}} \bigcup_{k=n-1}^{m} \mathbb{Q}^k_{-c} = Q^{n-1}_{-c}$$

and so $Q^{n-1}_{-c} \cap S_{-c}(h) \neq \emptyset$. Since $j^k_c(s_c)$ strongly believes Q^n_{-c}, it follows that $j^k_c(s_c)(Q^n_{-c}|S_{-c}(h)) = 1$. Now observe that, by construction,

$$\beta_c(\upsilon^k_c(s_c))(\bigcup_{l=n}^{k-1} \text{diag}^l_{-c}|S_{-c}(h) \times T_{-c}) = j^k_c(s_c)(Q^n_{-c}|S_{-c}(h)) = 1.$$

Since $\bigcup_{l=n}^{k-1} \text{diag}^l_{-c} \subseteq \bigcup_{l=n}^{m} Q^l_{-c}$ and, by the induction hypothesis, $R^n_{-c} = \bigcup_{l=n}^{m} \mathbb{Q}^l_{-c}$, it follows that $\beta_c(\upsilon^k_c(s_c))(R^n_{-c}|S_{-c}(h) \times T_{-c}) = 1$, as desired. ■

Proof of Proposition C.1. Immediate from Lemmata C.1-C.2. ■

C.2 Proof of Proposition C.2

Say a strategy s_c is **justifiable** if there exists some CPS p_c so that $s_c \in \mathbb{BR}[\text{p}_c]$. Proposition C.2 follows from the following lemma.

Lemma C.3. *Suppose that the game is generic and let $[s^*_{-c}] \subsetneq S_{-c}$. If s^*_c is justifiable, then there exists some CPS p_c so that $[s^*_c] = \mathbb{BR}[\text{p}_c]$ and p_c does not strongly believe $[s^*_{-c}]$.*

To show the lemma, it will be useful to begin with a number of preliminary results.

Lemma C.4. *Fix a CPS* $\mathrm{p}_c \in \mathcal{C}(S_{-c}, \mathcal{E}_c)$ *so that* $[s_c] = \mathbb{BR}[\mathrm{p}_c]$ *and some* $r_c \notin [s_c]$*. There exists some* $h \in H_c \cup \{\phi\}$ *so that* $s_c, r_c \in S_c(h)$ *and* $s_c(h) \neq r_c(h)$*. Moreover, for any such* h*,*

$$\sum_{s_{-c} \in S_{-c}(h)} [\pi_c(s_c, s_{-c}) - \pi_c(r_c, s_{-c})] p_c(s_{-c}|S_{-c}(h)) > 0.$$

Proof. Fix $[s_c] \subseteq \mathbb{BR}[\mathrm{p}_c]$ and $r_c \notin [s_c]$. Then, for all $h \in H_c \cup \{\phi\}$ with $s_c, r_c \in S_c(h)$

$$\sum_{s_{-c} \in S_{-c}(h)} [\pi_c(s_c, s_{-c}) - \pi_c(r_c, s_{-c})] p_c(s_{-c}|S_{-c}(h)) \geq 0. \tag{3}$$

Since $r_c \notin [s_c]$, there exists some $h^* \in H_c$ so that $s_c, r_c \in S_c(h^*)$ and $s_c(h^*) \neq r_c(h^*)$. We will suppose that Equation (3) holds with equality at $h = h^*$ and construct a new strategy r_c^* with $r_c^* \notin [s_c]$ and $r_c^* \in \mathbb{BR}[\mathrm{p}_c]$. This establishes the result.

Construct the strategy r_c^* as follows: First, for each information set h with either $S(h) \cap S(h^*) = \emptyset$ or $S(h) \subsetneq S(h^*)$, set $r_c^*(h) = s_c(h)$. Second, for each information set h with $S(h) \subseteq S(h^*)$ and $p_c(S_{-c}(h)|S_{-c}(h^*)) > 0$, set $r_c^*(h) = r_c(h)$. Finally, for all remaining information sets, choose r_c^* to satisfy the following condition: If $r_c^* \in S_c(h)$, then r_c^* solves

$$\max_{S_c(h)} \sum_{s_{-c} \in S_{-c}(h)} \pi_c(\cdot, s_{-c}) p_c(s_{-c}|S_{-c}(h)). \tag{4}$$

The fact that we can choose r_c^* in this way follows from Lemma A.3.[14]

Observe that $r_c^* \notin [s_c]$. Also observe that r_c^* is a best response under $p_c(\cdot|S_{-c}(h^*))$ given $S_c(h^*)$. To see this, fix some $s_{-c} \in \operatorname{Supp} p_c(\cdot|S_{-c}(h^*))$. Since $(s_c, s_{-c}), (r_c, s_{-c}) \in S_c(h^*) \times S_{-c}(h^*) = S(h^*)$, it follows from the construction that $(r_c^*, s_{-c}) \in S(h^*)$. Thus, $r_c^* \in S_c(h^*)$. Moreover, by construction, if $s_{-c} \in \operatorname{Supp} p_c(\cdot|S_{-c}(h^*))$ then $\zeta(r_c^*, s_{-c}) = \zeta(r_c, s_{-c})$. So, since r_c is a best response under $p_c(\cdot|S_{-c}(h^*))$ given $S_c(h^*)$, it follows that r_c^* is also is a best response under $p_c(\cdot|S_{-c}(h^*))$ given $S_c(h^*)$.

We will show that $r_c^* \in \mathbb{BR}[\mathrm{p}_c]$. Specifically, fix an information set $h \in H_c \backslash \{h^*\}$ with $r_c^* \in S_c(h)$. We will show that r_c^* is a best response under $p_c(\cdot|S_{-c}(h))$ given $S_c(h)$.

First, suppose that $S(h^*) \cap S(h) = \emptyset$. Fix some $p_c(s_{-c}|S_{-c}(h)) > 0$. By construction, $\zeta(r_c^*, s_{-c}) = \zeta(s_c, s_{-c})$. Since s_c is a best response under $p_c(\cdot|S_{-c}(h))$ given $S_c(h)$, it follows that r_c^* is also a best response under $p_c(\cdot|S_{-c}(h))$ given $S_c(h)$.

Second, suppose that $h \neq h^*$, $S(h) \subseteq S(h^*)$, and $p_c(S_{-c}(h)|S_{-c}(h^*)) > 0$. Since r_c^* is a best response under $p_c(\cdot|S_{-c}(h^*))$ given $S_c(h^*)$, Definition 2.2 and Lemma A.3 give that r_c^* is a best response under $p_c(\cdot|S_{-c}(h))$ given $S_c(h)$. Third, suppose that $h \neq h^*$, $S(h) \subseteq S(h^*)$, and $p_c(S_{-c}(h)|S_{-c}(h^*)) = 0$. In that case, by assumption, r_c^* is a best response under $p_c(\cdot|S_{-c}(h))$ given $S_c(h)$.

Finally, suppose that $S(h^*) \subsetneq S(h)$. Fix some $p_c(s_{-c}|S_{-c}(h)) > 0$. If $s_{-c} \notin S_{-c}(h^*)$, then $\zeta(r_c^*, s_{-c}) = \zeta(s_c, s_{-c})$. (This is by construction.) If $s_{-c} \in S_{-c}(h^*)$, then $\zeta(r_c^*, s_{-c}) = \zeta(r_c, s_{-c})$: Observe that $S_{-c}(h^*) \subseteq S_{-c}(h)$; so, by Definition 2.2, $p_c(s_{-c}|S_{-c}(h)) > 0$ implies $p_c(s_{-c}|S_{-c}(h^*)) > 0$. By construction, for any s_{-c} with $p_c(s_{-c}|S_{-c}(h^*)) > 0$, $\zeta(r_c^*, s_{-c}) = \zeta(r_c, s_{-c})$.

[14] Specifically: Let $\bar{H}_c^0$ be the set of all $h \in H_c$ with $S(h) \subseteq S(h^*)$, $p_c(S_{-c}(h)|S_{-c}(h^*)) = 0$, and $r_c \in S_c(h)$. Choose some $h^1 \in \bar{H}_c^0$ and note that $r_c^* \in S_c(h^1)$. Choose r_c^1 to solve Equation (4) for $h = h^1$ and set $r_c^*(h) = r_c^1(h)$. Then define $\bar{H}_c^1$ to be the set $h \in \bar{H}_c^0$ so that $r_c^1 \in S_c(h)$ and, if $S_{-c}(h) \subseteq S_{-c}(h^1)$, then $p_c(S_{-c}(h)|S_{-c}(h^1)) = 0$. Proceed inductively, until some $\bar{H}_c^K = \emptyset$ has been constructed. Then, "fill in" $r_c^*(h)$ arbitrarily at all information sets h for which it has not been defined. (Note, r_c^* precludes those information sets.)

Let $\alpha \equiv p_c(S_{-c}(h) \backslash S_{-c}(h^*) | S_{-c}(h)) > 0$. If $\alpha > 0$, let μ_c be $p_c(\cdot | S_{-c}(h))$ conditional on $S_{-c}(h) \backslash S_{-c}(h^*)$. If $\alpha = 0$, let $\mu_c 0$ be the zero measure. Then

$$\sum_{s_{-c} \in S_{-c}(h)} [\pi_c(s_c, s_{-c}) - \pi_c(r_c^*, s_{-c})] p_c(s_{-c} | S_{-c}(h)) =$$
$$\alpha \sum_{s_{-c} \in S_{-c}(h)} [\pi_c(s_c, s_{-c}) - \pi_c(r_c^*, s_{-c})] \mu(s_{-c}) + (1-\alpha) \sum_{s_{-c} \in S_{-c}(h)} [\pi_c(s_c, s_{-c}) - \pi_c(r_c^*, s_{-c})] p_c(s_{-c} | S_{-c}(h^*)).$$

Note that

$$\alpha \sum_{s_{-c} \in S_{-c}(h)} [\pi_c(s_c, s_{-c}) - \pi_c(r_c^*, s_{-c})] \mu(s_{-c}) = 0$$

since $\mu_c(s_{-c}) > 0$ implies $\zeta(s_c, s_{-c}) = \zeta(r_c^*, s_{-c})$. Also note that

$$(1-\alpha) \sum_{s_{-c} \in S_{-c}(h)} [\pi_c(s_c, s_{-c}) - \pi_c(r_c^*, s_{-c})] p_c(s_{-c} | S_{-c}(h^*)) = 0,$$

since both s_c and r_c^* are a best response under $p_c(\cdot | S_{-c}(h^*))$. Thus,

$$\sum_{s_{-c} \in S_{-c}(h)} [\pi_c(s_c, s_{-c}) - \pi_c(r_c^*, s_{-c})] p_c(s_{-c} | S_{-c}(h)) = 0.$$

Now, it follows from the fact that $s_c \in S_c(h^*) \subseteq S_c(h)$ is a best response under $p_c(\cdot | S_{-c}(h))$ given $S_c(h)$ that r_c is also a best response under $p_c(\cdot | S_{-c}(h))$ given $S_c(h)$. ■

Lemma C.5. *Fix some $h^* \in H_c \cup \{\phi\}$ so that $s_c^* \in S_c(h^*)$, $s_{-c}^* \notin S_{-c}(h^*)$ and, for all $h \in H_c \cup \{\phi\}$ with $S(h^*) \subsetneq S(h)$, $s_{-c}^* \in S_{-c}(h)$. Then $\zeta(s_c^*, s_{-c}^*) = \zeta(r_c, s_{-c}^*)$ impies $r_c \in S_c(h^*)$.*

Proof. We show the contrapositive. Suppose that $r_c \notin S_c(h^*)$. There exists some $(s_c^*, r_{-c}) \in S(h^*)$ so that $(r_c, r_{-c}) \notin S(h^*)$. Let n be the last common predecessor of $\zeta(s_c^*, r_{-c})$ and $\zeta(r_c, r_{-c})$. Note that there exists some $h \in H_c$ so that $n \in h$ and $s_c^*(h) \neq r_c(h)$. Observe that $S(h) \cap S(h^*) \neq \emptyset$. As such, either $S(h) \subseteq S(h^*)$ or $S(h^*) \subseteq S(h)$. Since $r_c \in S_c(h)$ but $r_c \notin S_c(h^*)$, it follows that $S(h^*) \subsetneq S(h)$. By construction, $s_{-c}^* \in S_{-c}(h)$. Thus, $\zeta(s_c^*, s_{-c}^*) \neq \zeta(r_c, s_{-c}^*)$. ■

Proof of Lemma C.3. Since the game is generic and s_c^* is justifiable, there exists some CPS p_c so that $[s_c^*] = \mathbb{BR}[\mathrm{p}_c]$. If p_c does not strongly believe $[s_{-c}^*]$, then we are done. So throughout we suppose otherwise. We will show that we can tilt p_c to construct a new CPS that satisfies the desired properties. We divide the argument into two cases.

Case A. Suppose that, for each $h \in H_c$ with $s_c^* \in S_c(h)$, $s_{-c}^* \in S_{-c}(h)$. So, for each $h \in H_c$ with $s_c^* \in S_c(h)$, $p_c(s_{-c}^* | S_{-c}(h)) = 1$. Lemma C.4 then implies that $\pi_c(s_c^*, s_{-c}^*) > \pi_c(s_c, s_{-c}^*)$ for all $s_c \in S_c \backslash [s_c^*]$.

Choose some $r_{-c}^* \in S_{-c} \backslash [s_{-c}^*]$. For each $\varepsilon \in (0, 1)$, construct a CPS q_c^ε so that

$$q_c^\varepsilon(s_{-c}^* | S_{-c}) = 1 - \varepsilon \qquad \text{and} \qquad q_c^\varepsilon(r_{-c}^* | S_{-c}) = \varepsilon$$

and, for each $h \in H_c$ with $S_{-c}(h) \cap \{s_{-c}^*, r_{-c}^*\} = \emptyset$, $q_c^\varepsilon(\cdot | S_{-c}(h)) = p_c(\cdot | S_{-c}(h))$. Note, the unique CPS q_c^ε that satisfies these conditions does not strongly believe $[s_{-c}^*]$.

Now observe that we can find some $\bar{\varepsilon} > 0$ so that for each $\varepsilon \in (0, \bar{\varepsilon})$ the following holds: If $h \in H_c$ with $s_c^* \in S_c(h)$, then

$$\sum_{s_{-c} \in S_{-c}} [\pi_c(s_c^*, s_{-c}) - \pi_c(r_c, s_{-c})] q_c^{\varepsilon}(s_{-c}|S_{-c}) =$$
$$(1-\varepsilon)[\pi_c(s_c^*, s_{-c}^*) - \pi_c(s_c, s_{-c}^*)] + \varepsilon[\pi_c(s_c^*, r_{-c}^*) - \pi_c(s_c, r_{-c}^*)] > 0$$

for each $r_c \in S_c(h)$. Thus, $\mathbb{BR}[\mathrm{q}_c^{\varepsilon}] = [s_c^*]$ for all $\varepsilon \in (0, \bar{\varepsilon})$.

Case B. Suppose that there exists some $h^* \in H_c$ so that $s_c^* \in S_c(h^*)$ but $s_{-c}^* \notin S_{-c}(h^*)$. Choose h^* so that, if $S(h^*) \subsetneq S(h)$, then $s_{-c}^* \in S_{-c}(h)$. Let $\mu_c^* = p_c(\cdot|S_{-c}(h^*))$ and observe that $\mu_c^*([s_{-c}^*]) = 0$ since $s_{-c}^* \notin S_{-c}(h^*)$. For each $\varepsilon \in (0,1)$, construct a CPS $\mathrm{q}_c^{\varepsilon}$ so that

$$q_c^{\varepsilon}(s_{-c}|S_{-c}) = \begin{cases} 1-\varepsilon & \text{if } s_{-c} = s_{-c}^* \\ \varepsilon\mu_c^*(s_{-c}) & \text{if } s_{-c} \neq s_{-c}^*. \end{cases}$$

and, for each $h \in H_c$, with $S_{-c} \cap (\{s_{-c}^*\} \cup \operatorname{Supp} \mu_c^*) = \emptyset$, $q_c^{\varepsilon}(\cdot|S_{-c}(h)) = p_c(\cdot|S_{-c}(h))$. Note, the unique CPS $\mathrm{q}_c^{\varepsilon}$ that satisfies these conditions does not strongly believe $[s_{-c}^*]$.

Step 1: We begin by showing that, for each $r_c \in S_c$, there exists some $\bar{\varepsilon}(r_c) > 0$ so that the following holds: For all $\varepsilon \in (0, \bar{\varepsilon}(r_c))$,

$$\sum_{s_{-c} \in S_{-c}} [\pi_c(s_c^*, s_{-c}) - \pi_c(r_c, s_{-c})] q_c^{\varepsilon}(s_{-c}|S_{-c}) \begin{cases} > 0 & \text{if } \zeta(r_c, s_{-c}^*) \neq \zeta(s_c^*, s_{-c}^*) \\ \geq 0 & \text{if } \zeta(r_c, s_{-c}^*) = \zeta(s_c^*, s_{-c}^*). \end{cases} \tag{5}$$

First, suppose that $\zeta(r_c, s_{-c}^*) \neq \zeta(s_c^*, s_{-c}^*)$. Then, there exists some $\tilde{h}$ so that $(s_c^*, s_{-c}^*), (r_c, s_{-c}^*) \in S(\tilde{h})$ and $s_c^*(\tilde{h}) \neq r_c(\tilde{h})$. Moreover, $p_c(s_{-c}^*|S_{-c}(\tilde{h})) = 1$. Thus, applying Lemma C.4, $\pi_c(s_c^*, s_{-c}^*) > \pi_c(r_c, s_{-c}^*)$. It follows that there exists some $\bar{\varepsilon}(r_c) > 0$ so that, for all $\varepsilon \in (0, \bar{\varepsilon}(r_c))$,

$$\sum_{s_{-c} \in S_{-c}} [\pi_c(s_c^*, s_{-c}) - \pi_c(r_c, s_{-c})] q_c^{\varepsilon}(s_{-c}|S_{-c}) =$$
$$(1-\varepsilon)[\pi_c(s_c^*, s_{-c}^*) - \pi_c(r_c, s_{-c}^*)] +$$
$$\varepsilon \sum_{s_{-c} \in S_{-c}} [\pi_c(s_c^*, s_{-c}) - \pi_c(r_c, s_{-c})] \mu_c^*(s_{-c}) > 0$$

Second, suppose that $\zeta(r_c, s_{-c}^*) = \zeta(s_c^*, s_{-c}^*)$. In this case, $\pi_c(s_c^*, s_{-c}^*) - \pi_c(r_c, s_{-c}^*) = 0$. Moreover, if $s_c^* \in S_c(h^*)$, then $r_c \in S_c(h^*)$. (See Lemma C.5.) Since s_c^* is a best response under μ_c^* given $S_c(h^*)$, it follows that

$$\sum_{s_{-c} \in S_{-c}} [\pi_c(s_c^*, s_{-c}) - \pi_c(r_c, s_{-c})] \mu_c^*(s_{-c}) \geq 0.$$

As such,

$$\sum_{s_{-c}\in S_{-c}} [\pi_c(s_c^*, s_{-c}) - \pi_c(r_c, s_{-c})] q_c^{\varepsilon}(s_{-c}|S_{-c}) =$$
$$(1-\varepsilon)[\pi_c(s_c^*, s_{-c}^*) - \pi_c(r_c, s_{-c}^*)] +$$
$$\varepsilon \sum_{s_{-c}\in S_{-c}} [\pi_c(s_c^*, s_{-c}) - \pi_c(r_c, s_{-c})]\mu_c^*(s_{-c}) \geq 0$$

for all $\varepsilon > 0$.

Step 2: Take $\bar{\varepsilon} = \min\{\bar{\varepsilon}(r_c) : r_c \in S_c\}$. We will show that $[s_c^*] \subseteq \mathbb{BR}[\mathrm{q}_c^{\varepsilon}]$ for all $\varepsilon \in (0, \bar{\varepsilon})$. To do so, begin by noting that Equation (5) holds for all $r_c \in S_c$, provided $\varepsilon \in (0, \bar{\varepsilon})$. To complete the argument, it suffices to show that, if $h \in H_c$ with $s_c \in S_c(h)$ then either: $q_c^{\varepsilon}(\cdot|S_{-c}(h)) = q_c^{\varepsilon}(\cdot|S_{-c})$ or $q_c^{\varepsilon}(\cdot|S_{-c}(h)) = p_c(\cdot|S_{-c}(h))$. From this the conclusion will follow.

First, suppose that $S(h^*) \subseteq S(h)$. In that case, $q_c^{\varepsilon}(\cdot|S_{-c}(h)) = q_c^{\varepsilon}(\cdot|S_{-c})$. Second, suppose that $S(h) \subsetneq S(h^*)$. In that case, $q_c^{\varepsilon}(\cdot|S_{-c}(h)) = p_c(\cdot|S_{-c}(h))$. Finally, suppose that $S(h^*) \cap S(h) = \emptyset$. In that case, $s_c \in S(h^*) \cap S(h)$ and, so, $S_{-c}(h^*) \cap S_{-c}(h) = \emptyset$. (See Lemma A.2.) From this, $q_c^{\varepsilon}(\mathrm{Supp}\, \mu_c^*|S_{-c}(h)) = 0$ and so $q_c^{\varepsilon}(\cdot|S_{-c}(h)) = p_c(\cdot|S_{-c}(h))$.

Step 3: We now show that, for all $\varepsilon \in (0, \bar{\varepsilon})$, $\mathbb{BR}[\mathrm{q}_c^{\varepsilon}] \subseteq [s_c^*]$. To see this, fix some $r_c \notin [s_c]$. Then there exists some $h \in H_c \cup \{\phi\}$ so that $s_c, r_c \in S_c(h)$ and

$$\sum_{s_{-c}\in S_{-c}} [\pi_c(s_c^*, s_{-c}) - \pi_c(r_c, s_{-c})] p_c(s_{-c}|S_{-c}(h)) > 0.$$

(See Lemma C.4.) If $q_c^{\varepsilon}(\cdot|S_{-c}(h)) = p_c(\cdot|S_{-c}(h))$, then certainly $r_c \notin \mathbb{BR}[\mathrm{q}_c^{\varepsilon}]$. If $q_c^{\varepsilon}(\cdot|S_{-c}(h)) \neq p_c(\cdot|S_{-c}(h))$, then $S(h^*) \subsetneq S(h)$. In that case,

$$\sum_{s_{-c}\in S_{-c}} [\pi_c(s_c^*, s_{-c}) - \pi_c(r_c, s_{-c})] p_c(s_{-c}|S_{-c}(h)) = \pi_c(s_c^*, s_{-c}^*) - \pi_c(r_c, s_{-c}^*) > 0.$$

Thus, $\zeta(s_c^*, s_{-c}^*) \neq \zeta(r_c, s_{-c}^*)$ and, so, by Equation (5) $r_c \notin \mathbb{BR}[\mathrm{q}_c^{\varepsilon}]$. ■

Appendix D Generic Games

The first half of the appendix focuses on no relevant convexities. We then turn to the condition of no relevant ties.

D.1 No Relevant Convexities

In this appendix, we focus on a weaker version of NRC:

Definition D.1. The game satisfies **weak no relevant convexities** (**WNRC**) if, for each $h \in H_c$, the following holds: If $s_c, r_c \in S_c(h)$ and r_c supports s_c with respect to some $S_c(h) \times (X_{-c} \cap S_{-c}(h))$, then $\zeta(s_c, s_{-c}) = \zeta(r_c, s_{-c})$ for each $s_{-c} \in X_{-c} \cap S_{-c}(h)$.

NRC requires the convexity condition hold for all subsets X_{-c} in the strategic-form. It follows that it also holds at each information set, as required by WNRC. Thus, if a game satisfies NRC, then it also satisfies WNRC. We will show:

Proposition D.1. *Fix a game that satisfies WNRC. Then $(Q^0, \ldots, Q^m)$ satisfies the m-BRP if and only if Q^1 is non-empty and, for each $n = 1, \ldots, m$ and each $s_c \in Q_c^n$,*

(i) there exists an array $\mathrm{p}_c \in \mathcal{A}(S_{-c}, \mathcal{E}_c)$ that strongly believes $Q^0_{-c}, \ldots, Q^{n-1}_{-c}$ with $s_c \in \mathbb{BR}[\mathrm{p}_c]$, and

(ii) $[s_c] \subseteq Q_c^n$.

Remark D.1. The proof shows that Proposition D.1 obtains if the following weaker requirement is met: Fix some $h \in H_c$ and some $s_c \in S_c(h)$. Suppose that there exists some $\mu_c \in \mathcal{P}(S_{-c}(h))$ so that s_c is a best response under μ_c given $S_c(h)$. Then, $r_c \in S_c(h)$ supports s_c with respect to $S_c(h) \times \operatorname{Supp} \mu_c$ only if $\zeta(s_c, s_{-c}) = \zeta(r_c, s_{-c})$ for each $s_{-c} \in \operatorname{Supp} \mu_c$. □

To show this, it suffices to show the following:

Lemma D.1. *Fix a game that satisfies WNRC. Let $(Q^0, \ldots, Q^{m-1}, Q^m)$ be a decreasing sequence of strategy profiles. Suppose that, for each $s_c \in Q_c^m$,*

(i) there exists an array $\mathrm{p}_c \in \mathcal{A}(S_{-c}, \mathcal{E}_c)$ that strongly believes $Q^0_{-c}, \ldots, Q^{m-1}_{-c}$ with $s_c \in \mathbb{BR}[\mathrm{p}_c]$ and

(ii) $[s_c] \subseteq Q_c^n$.

Then there exists a CPS $q_c \in \mathcal{C}(S_{-c}, \mathcal{E}_c)$ so that

(a) $s_c \in \mathbb{BR}[q_c]$,

(b) q_c strongly believes $Q^0_{-c}, \ldots, Q^{m-1}_{-c}$, and

(c) $\mathbb{BR}[q_c] = [s_c]$.

To show this result, we will take two steps. First, we will show that, if the premise of the Lemma holds, then we can construct an array satisfying conditions (a)-(b) and a variant of (c). Second, we will show that, if we have an array satisfying these conditions, then we can construct a CPS satisfying conditions (a)-(b)-(c) of the Lemma.

Constructing the Array The following Lemma will of use. (It follows from Lemmata D.2-D.3-D.4 in Brandenburger, Friedenberg and Keisler, 2008.)

Lemma D.2. *Suppose that s_c is optimal under $\mu_c \in \mathcal{P}(S_{-c})$ amongst strategies in $X_c \subseteq S_c$. Then there exists some $\nu_c \in \mathcal{P}(S_{-c})$ with $\operatorname{Supp} \nu_c = \operatorname{Supp} \mu_c$ so that the following holds: r_c is a best response under ν_c given X_c if and only if r_c supports s_c with respect to $X_c \times \operatorname{Supp} \mu_c$.*

Lemma D.3. *Fix a game that satisfies WNRC. Suppose that s_c is sequentially optimal under the array $\mathrm{p}_c = (p_c(\cdot|S_{-c}(h)) : h \in H_c \cup \{\phi\})$. Then there exists an array $\mathrm{q}_c = (q_c(\cdot|S_{-c}(h)) : h \in H_c \cup \{\phi\})$ so that the following hold:*

(i) For each $h \in H_c$ with $s_c \in S_c(h)$, if $r_c \in S_c(h)$ is a best response under $q_c(\cdot|S_{-c}(h))$ given $S_c(h)$, then $\zeta(s_c, s_{-c}) = \zeta(r_c, s_{-c})$ for all $s_{-c} \in \operatorname{Supp} q_c(\cdot|S_{-c}(h))$.

(ii) Each $r_c \in [s_c]$ *is a sequential best reply under* q_c*.*

(iii) The array q_c *strongly believes* E_{-c} *if and only if the array* p_c *strongly believes* E_{-c}*.*

Proof. By Lemma D.2, we can construct an array q_c so that, for each $h \in H_c \cup \{\phi\}$, the following hold:

(a) $\operatorname{Supp} q_c(\cdot|S_{-c}(h)) = \operatorname{Supp} p_c(\cdot|S_{-c}(h))$, and

(b) if $s_c \in S_c(h)$, then $r_c \in S_c(h)$ is a best response under $q_c(\cdot|S_{-c}(h))$ given $S_c(h)$ if and only if r_c supports s_c with respect to $S_c(h) \times \operatorname{Supp} p_c(\cdot|S_{-c}(h))$.

Observe that (a)-(b) and WNRC gives part (i), from which (ii) follows. Moreover, (a) gives part (iii). ■

Corollary D.1. *Fix a game that satisfies WNRC. Let* $(Q^0, \ldots, Q^{m-1}, Q^m)$ *be a decreasing sequence of strategy profiles. Suppose that, for each* $s_c \in Q_c^m$*, there exists an array* $\mathrm{p}_c \in \mathcal{A}(S_{-c}, \mathcal{E}_c)$ *so that*

(i) $s_c \in \mathbb{BR}[\mathrm{p}_c]$

(ii) p_c *strongly believes* $Q_{-c}^0, \ldots, Q_{-c}^{m-1}$*, and*

(iii) $[s_c] \subseteq Q_c^n$*.*

Then there exists an array $q_c \in \mathcal{A}(S_{-c}, \mathcal{E}_c)$ *so that*

(a) $s_c \in \mathbb{BR}[q_c]$*,*

(b) q_c *strongly believes* $Q_{-c}^0, \ldots, Q_{-c}^{m-1}$*, and*

(c) for each $h \in H_c$ *with* $s_c \in S_c(h)$*, if* $r_c \in S_c(h)$ *is a best response under* $q_c(\cdot|S_{-c}(h))$ *given* $S_c(h)$*, then* $\zeta(s_c, s_{-c}) = \zeta(r_c, s_{-c})$ *for all* $s_{-c} \in \operatorname{Supp} q_c(\cdot|S_{-c}(h))$*.*

Constructing the CPS In what follows, we will fix an array $\mathrm{p}_c \in \mathcal{A}(S_{-c}, \mathcal{E}_c)$ with $s_c \in \mathbb{BR}[\mathrm{p}_c]$ that satisfies the following property:

> Property [*]: For each $h \in H_c$ with $s_c \in S_c(h)$, if $r_c \in S_c(h)$ is a best response under $p_c(\cdot|S_{-c}(h))$ then $\zeta(s_c, s_{-c}) = \zeta(r_c, s_{-c})$ for all $s_{-c} \in \operatorname{Supp} p_c(\cdot|S_{-c}(h))$.

Lemma D.4. *Fix a game that satisfies WNRC. Suppose there is an array* $\mathrm{p}_c \in \mathcal{A}(S_{-c}, \mathcal{E}_c)$ *with* $s_c \in \mathbb{BR}[\mathrm{p}_c]$ *that satisfies Property [*]. Then there exists a CPS* $\mathrm{q}_c \in \mathcal{C}(S_{-c}, \mathcal{E}_c)$ *so that the following hold:*

(i) $\mathbb{BR}[\mathrm{q}_c] = [s_c]$*, and*

(ii) if p_c *strongly believes* E_{-c}*, then* q_c *strongly believes* E_{-c}*.*

Fix some array $\mathrm{p}_c \in \mathcal{A}(S_{-c}, \mathcal{E}_c)$ satisfying Property [*]. We construct a CPS $\mathrm{q}_c = (q_c(\cdot|S_{-c}(h)) : h \in H_c \cup \{\phi\})$ inductively: Let $H_c^0 = H_c \cup \{\phi\}$ and choose $h^0 = \phi \in H_c^0$ and observe that $S_{-c}(\phi) = S_{-c}$. Set $q_c(\cdot|S_{-c}) = p_c(\cdot|S_{-c})$. Define $\overline{H}_c^0$ to be the set of $h \in H_c$ so that $S_{-c}(h) \subseteq S_{-c}$ and $q_c(S_{-c}(h)|S_{-c}) > 0$. For each $h \in \overline{H}_0^c$, set

$$q_c(s_{-c}|S_{-c}(h)) = \frac{q_c(s_{-c}|S_{-c})}{q_c(S_{-c}(h)|S_{-c})}$$

for all $s_{-c} \in S_{-c}(h)$. Note, $h^0 \in \overline{H}_c^0$.

Assume the sets H_c^k and $\overline{H}_c^k$ have been defined. Set $H_c^{k+1} = H_c^k \backslash \overline{H}_c^k$. If $H_c^{k+1} = \emptyset$, then we are done. If not, choose some $h^{k+1} \in H_c^{k+1}$ that satisfies the following requirements:

(i) Either $s_c \in S_c(h^{k+1})$ or, for all $h \in H_c^{k+1}$, $s_c \notin S_c(h)$.

(ii) There is no $h \in H_c^{k+1}$ so that $S_{-c}(h^{k+1}) \subsetneq S_{-c}(h)$.

(iii) If $S_{-c}(h^{k+1}) = S_{-c}(h)$, then either $S_c(h) \subseteq S_c(h^{k+1})$ or $S_c(h) \cap S_c(h^{k+1}) = \emptyset$.

Set $q_c(\cdot|S_{-c}(h^{k+1})) = p_c(\cdot|S_{-c}(h^{k+1}))$. Define $\overline{H}_c^{k+1}$ to be the set of h so that $S_{-c}(h) \subseteq S_{-c}(h^{k+1})$ and $q_c(S_{-c}(h)|S_{-c}(h^{k+1})) > 0$. For each $h \in \overline{H}_{k+1}^c$, set

$$q_c(s_{-c}|S_{-c}(h)) = \frac{q_c(s_{-c}|S_{-c})}{q_c(S_{-c}(h)|S_{-c}(h^{k+1}))}$$

for all $s_{-c} \in S_{-c}(h)$.

It might be useful to recap the construction: We begin by identifying information sets $h^0, h^1, \ldots, h^K$. In keeping with the terminology in Siniscalchi (2016), we refer to these as *basic information sets*. We set $q_c(\cdot|S_{-c}(h^k))$ to coincide with the original CPS $p_c(\cdot|h^k)$. For any non-basic information set h, there is exactly one basic information h^k so that $S_{-c}(h) \subseteq S_{-c}(h^k)$ and $q_c(S_{-c}(h)|S_{-c}(h^k)) > 0$. Thus we construct the belief $q_c(\cdot|S_{-c}(h))$ from $q_c(\cdot|S_{-c}(h^k))$ by conditioning on $S_{-c}(h)$. The construction obviously yields a CPS. We note the following:

Lemma D.5. *If $h \in \overline{H}_c^k$ and $s_c \in S_c(h)$, then $S(h) \subseteq S(h^k)$.*

Proof. Suppose, contra hypothesis, that $S(h)$ is not contained in $S(h^k)$. Then either $S(h^k) \subsetneq S(h)$ or $S(h) \cap S(h^k) = \emptyset$. The first of these cannot happen by construction. So, it must be that $S(h) \cap S(h^k) = \emptyset$. Since $h \in \overline{H}_c^k$, $S_{-c}(h) \cap S_{-c}(h^k) \neq \emptyset$; it follows that $S_c(h) \cap S_c(h^k) = \emptyset$. (See Lemma A.2.) Since $s_c \in S_c(h)$, it follows that $s_c \notin S_c(h^k)$. But this contradicts how we constructed the basic information sets h^k. ■

We prove Lemma D.4 by showing that (i) $\mathbb{BR}[\mathrm{q}_c] = [s_c]$, and (ii) if p_c strongly believes E_{-c}, then q_c strongly believes E_{-c}.

Lemma D.6. $\mathbb{BR}[\mathrm{q}_c] = [s_c]$.

Proof. First we show that $[s_c] \subseteq \mathbb{BR}[\mathrm{q}_c]$. To do so, it suffices to show that $s_c \in \mathbb{BR}[\mathrm{q}_c]$. Toward that end, fix some $h \in H_c$ with $s_c \in S_c(h)$. Observe that there exist a k so that $h \in \overline{H}_c^k$, i.e., there exists a basic h^k so that $q_c(\cdot|S_{-c}(h))$ is derived from $p_c(\cdot|S_{-c}(h^k))$ by conditioning. (Note, h may well be h^k.) By construction, s_c is optimal under $q_c(\cdot|S_{-c}(h^k))$ given all strategies in $S_c(h^k)$. It follows from Lemmata D.5-A.3 that s_c is a best response under $q_c(\cdot|S_{-c}(h))$ given all strategies in $S_c(h)$.

Next fix some $r_c \in \mathbb{BR}[\mathrm{q}_c]$ and suppose that $r_c \notin [s_c]$. Then there is an information set $h \in H_c$ so that $s_c, r_c \in S_c(h)$ and $s_c(h) \neq r_c(h)$. Suppose that r_c is a best response under $q_c(\cdot|S_c(h))$ given $S_c(h)$. Note, there exists k such that $h \in \overline{H}_c^k$ and, by Lemma D.5, $S(h) \subseteq S(h^k)$. Using Lemma A.3, r_c is a best response under $q_c(\cdot|S_{-c}(h^k)) = p_c(\cdot|S_{-c}(h^k))$ given $S_c(h^k)$. But now observe that there exists some $s_{-c} \in \operatorname{Supp} p_c(\cdot|S_{-c}(h^k)) \cap S_{-c}(h)$ with $\zeta(s_c, s_{-c}) \neq \zeta(r_c, s_{-c})$. This contradicts Property [*]. ■

Lemma D.7. *If p_c strongly believes E_{-c}, then q_c strongly believes E_{-c}.*

Proof. Fix an information set $h \in H_c$ so that $E_{-c} \cap S_{-c}(h) \neq \emptyset$. There exists some $h^k \in H_c$ so that $S_{-c}(h) \subseteq S_{-c}(h^k)$, $p_c(S_{-c}(h)|S_{-c}(h^k)) > 0$ and, for every $s_{-c} \in S_{-c}(h)$,

$$q_c(s_{-c}|S_{-c}(h)) = \frac{p_c(s_{-c}|S_{-c}(h^k))}{p_c(S_{-c}(h)|S_{-c}(h^k))}.$$

Since $S_{-c}(h) \subseteq S_{-c}(h^k)$, $E_{-c} \cap S_{-c}(h^k) \neq \emptyset$. If p_c strongly believes E_{-c} then $p(E_{-c}|S_{-c}(h^k)) = 1$ and so $q(E_{-c}|S_{-c}(h)) = 1$. ■

Recap Let us sum up. Proposition D.1 (resp. Proposition 6.1) follows immediately from Lemma D.1. In turn, Lemma D.1 follows from Corollary D.1 and Lemma D.4. Moreover, as an implication of Lemma D.1, we have the following corollary.

Corollary D.2. *If a game satisfies WNRC, then it is generic.*

Proof. Fix some $s_c \in \mathbb{BR}[\mathrm{p}_c]$ for some $\mathrm{p}_c \in \mathcal{C}(S_{-c}, \mathcal{E}_c)$. Then there exists a 1-BRP (Q^0, Q^1) such that $[s_c] \in Q^1_c$. By Lemma D.1, it follows that there exists some $\mathrm{q}_c \in \mathcal{C}(S_{-c}, \mathcal{E}_c)$ with $[s_c] = \mathbb{BR}[\mathrm{q}_c]$. Thus the game is generic. ■

One implication of Proposition D.1 (resp. Proposition 6.1) is that, when WNRC (resp. NRC) is satisfied, we can forgo using CPS's and focus on arrays. This would not be the case absent WNRC (resp. NRC). The central difficulty comes from condition (BRP.3) of the m-BRP. Specifically, begin with a decreasing sequence of strategy profiles $(Q^0, \ldots, Q^{m-1}, Q^m)$. In addition, suppose that $s_c \in Q^m_c$ so that, for some array $\mathrm{p}_c \in \mathcal{A}(S_{-c}, \mathcal{E}_c)$, conditions (BRP.1)-(BRP.2)-(BRP.3) are satisfied. We can use this array to construct a CPS $\mathrm{q}_c \in \mathcal{C}(S_{-c}, \mathcal{E}_c)$ so that conditions (BRP.1)-(BRP.2) are satisfied. (This argument is standard.) But, condition (BRP.3) may fail for the constructed CPS. The next example makes this point.

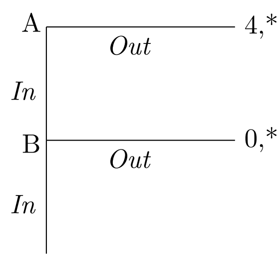

		Bob	
		L	*R*
Ann	*U*	4,*	2,*
	D	4,*	1,*

Figure D.1: Arrays Do Not Suffice

Example D.1. Consider the game in Figure D.1, which fails WNRC. Write h for the information set at which the simultaneous move game is played. Let p_a be an array so that $p_a(I\text{-}L|S_b) = 1$ and $p_a(I\text{-}L|S_b(h)) = p_a(I\text{-}R|S_b(h)) = \frac{1}{2}$. Observe that $\mathbb{BR}_a[\mathrm{p}_a] = \{O, I\text{-}U\}$.

We can use this array to construct a CPS $\mathrm{q}_a \in \mathcal{C}(S_b, \mathcal{E}_a)$: We set $q_a(I\text{-}L|S_b) = p_a(I\text{-}L|S_b) = 1$ and $q_a(I\text{-}L|S_b(h)) = q_a(I\text{-}L|S_b) = 1$. However, $\mathbb{BR}_a[\mathrm{q}_a] = \{O, I\text{-}U, I\text{-}D\}$, i.e., it contains an additional strategy. In fact, there is no CPS $\hat{\mathrm{q}}_a \in \mathcal{C}(S_b, \mathcal{E}_a)$ with $\mathbb{BR}_a[\hat{\mathrm{q}}_a] = \mathbb{BR}_a[\mathrm{p}_a]$. □

D.2 No Relevant Ties

A perfect information game can satisfy no relevant ties, even if it fails no relevant convexities. Nonetheless, if a perfect information game satisfies no relevant ties, it is generic.

Lemma D.8. *A perfect-information game satisfying no relevant ties is generic.*

Perfect-information is important for the result. The next example highlights this fact.

Example D.2. The game in Figure D.2 satisfies no relevant ties. Yet it is not generic: *Out* is optimal under a CPS p_a if and only if $p_a(L|S_b) = p_a(R|S_b) = \frac{1}{2}$. Thus, $\mathbb{BR}[\mathrm{p}_a] = \{Out, U, M\}$ and there is no CPS q_a with $\mathbb{BR}[\mathrm{q}_a] = [Out\,]$. □

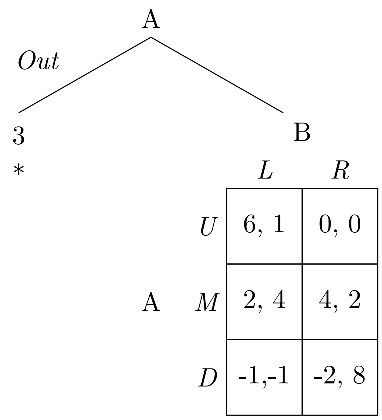

Figure D.2: No Relevant Ties

To show this lemma, we will need an auxiliary definition and result.

Definition D.2. Given a conditional probability space $(\Omega, \mathcal{E})$, call a CPS $p \in \mathcal{C}(\Omega, \mathcal{E})$ **degenerate** if, for each conditioning event E, there exists some $\omega \in E$ with $p(\omega|E) = 1$.

The following Lemma follows almost immediately from Ben-Porath (1997, Lemma 1.2.1).

Lemma D.9. *Fix a perfect-information game satisfying no relevant ties. If s_c is justifiable, then there exists some degenerate CPS $\mathrm{p}_c \in \mathcal{C}(S_{-c}, \mathcal{E}_c)$ so that $s_c \in \mathbb{BR}[\mathrm{p}_c]$.*

In a perfect-information game, we can identify an information set h with the unique node (or vertex) it contains. In that case, we will say an information set h precedes an information set h' if $h = \{v\}$, $h' = \{v'\}$, and v precedes v'. We will say that h strictly precedes h' if h precedes h' and $h \neq h'$. We will say that h weakly precedes h' if $h = h'$.

Proof of Lemma D.9. Let s_c be a justifiable strategy. Then, by Lemma 1.2.1 in Ben-Porath (1997), for each $S_{-c}(h) \in \mathcal{E}_c$ with $s_c \in S_c(h)$, we can find some $s^h_{-c} \in S_{-c}(h)$ so that $\pi_c(s_c, s^h_{-c}) \geq \pi_c(r_c, s^h_{-c})$ for all $r_c \in S_c(h)$. Use the collection $(s^h_{-c} : h \in H_c \cup \{\phi\})$ to form a CPS p_c.

We will inductively define the measures $p_c(\cdot|S_{-c}(h))$. For each $S_{-c}(h)$ with $s^{\{\phi\}}_{-c} \in S_{-c}(h)$, set $p_c(s^{\{\phi\}}_{-c}|S_{-c}(h)) = 1$. Next, fix an information set $h^* \in H_c$ where $p_c(\cdot|S_{-c}(h))$ has been defined for each h that strictly precedes h^* but for which $p_c(\cdot|S_{-c}(h^*))$ has not been defined. Set $p_c(s^{h^*}_{-c}|S_{-c}(h)) = 1$ for each $S_{-c}(h)$ with $s^{h^*}_{-c} \in S_{-c}(h)$. Proceeding along these lines, we define $p_c(\cdot|S_{-c}(h))$ for each conditioning event $S_{-c}(h)$.

It can be verified that, so defined, p_c is a CPS. Moreover, s_c is optimal under p_c: Given an information set $h \in H_c$ with $s_c \in S_c(h)$, there exists an information set h^* that precedes (perhaps weekly) h so that

$s_{-c} \in S_{-c}(h)$ and $p_c(s^{h^*}_{-c}|S_{-c}(h)) = 1$. Then, the claim follows from the fact that $S_c(h) \subseteq S_c(h^*)$ and the fact that $\pi_c(s_c, s^h_{-c}) \geq \pi_c(r_c, s^h_{-c})$ for all $r_c \in S_c(h)$. ■

Proof of Lemma D.8. Fix a perfect-information game satisfying no relevant ties and some strategy s_c that is justifiable. Then there exists some degenerate CPS $\mathrm{p}_c \in \mathcal{C}(S_{-c}, \mathcal{E}_c)$ so that $s_c \in \mathbb{BR}[\mathrm{p}_c]$. We will show that, if $r_c \notin [s_c]$, then $r_c \notin \mathbb{BR}[\mathrm{p}_c]$.

Fix some $r_c \notin [s_c]$. Then there exists some $h \in H_c$ with $s_c, r_c \in S_c(h)$ and $s_c(h) \neq r_c(h)$. Let $s_{-c} \in S_{-c}(h)$ with $\mathrm{p}_c(s_{-c}|S_{-c}(h)) = 1$. Since s_c is sequentially optimal under p_c, $\pi_c(s_c, s_{-c}) \geq \pi_c(r_c, s_{-c})$. But, since $\zeta(s_c, s_{-c}) \neq \zeta(r_c, s_{-c})$, no relevant ties implies $\pi_c(s_c, s_{-c}) > \pi_c(r_c, s_{-c})$. Thus, $r_c \notin \mathbb{BR}[\mathrm{p}_c]$. ■

Remark D.2. The proof of Lemma D.9 constructs a CPS p_c for which the property of Remark D.1 holds. Thus, a perfect information game satisfying NRT satisfies the weaker condition discussed in Remark D.1. This is an alternate route to show that the game is generic. □

Appendix E Centipede

Throughout this Appendix, fix an m-BRP $(Q^0, Q^1, \ldots, Q^m)$ of the centipede game. We will show that $Q^m_a \subseteq \mathrm{EFR}^m_a$. We begin with the following observation:

Observation E.1. *Observe that* $[\mathrm{in}]_\ell \cap Q^1_\ell = \emptyset$ *and so* $Q^1_\ell \times Q^1_{-\ell} \subseteq \mathrm{EFR}^1_\ell \times \mathrm{EFR}^1_{-\ell}$.

Lemma E.1. *One of the following must hold:*

(i) $[\mathrm{in}]_{-\ell} \cap Q^2_{-\ell} = \emptyset$, *or*

(ii) $[\mathrm{out}, \mathrm{V}_\ell] \cap Q^1_\ell = \emptyset$ *and* $V = 3$.

Proof. First, suppose that $[\mathrm{out}, \mathrm{V}]_\ell \subseteq Q^1_\ell$. In that case, any CPS strongly believes Q^1_ℓ must assign probability one to $[\mathrm{out}, \mathrm{V}_\ell]$ at node $V - 1$. (This uses Observation E.1, i.e., the fact that $[\mathrm{in}]_\ell \cap Q^1_\ell = \emptyset$.) Thus, $[\mathrm{in}]_{-\ell}$ is not a best response at node $V - 1$. From this $[\mathrm{in}]_{-\ell} \cap Q^2_{-\ell} = \emptyset$.

Second, suppose that $[\mathrm{out}, \mathrm{V}]_\ell \cap Q^1_\ell = \emptyset$. Let $\mathrm{p}_{-\ell}$ be a CPS that strongly believes Q^1_ℓ and note that $p_{-\ell}(\cdot|S_\ell)$ must assign probability one to

$$\{s_\ell : s_\ell(v) = \mathrm{out}_v \text{ for some } v \leq V - 2\}.$$

(That is, *ex ante*, $\mathrm{p}_{-\ell}$ assigns probability one to the game ending at some node $v \leq V - 2$, independent of the strategy that $-\ell$ plays.) If $V \geq 4$, then there is some node $\tilde{v} \leq V - 3$ at which $-\ell$ moves and $p_{-\ell}([\mathrm{out}, \tilde{v} + 1]_\ell | S_\ell(\tilde{v})) = 1$. Thus, at node $\tilde{v}$, $[\mathrm{out}, \tilde{v}]_{-\ell}$ is a unique best response. So certainly $[\mathrm{in}]_{-\ell} \cap Q^2_{-\ell} = \emptyset$. ■

Lemma E.2. *Fix some* $m = 3, \ldots, V$. *If* m *is odd then either*

(i) $[\mathrm{out}, V + 3 - m]_\ell \cap Q^m_\ell = \emptyset$, *or*

(ii) $[\mathrm{out}, \mathrm{V} + 2 - \mathrm{m}]_{-\ell} \cap Q^{m-1}_{-\ell} = \emptyset$ *and* $V \leq m + 1$.

And, if m *is even then either*

(i) $[\mathrm{out}, V + 3 - m]_{-\ell} \cap Q^m_{-\ell} = \emptyset$, *or*

(ii) $[\mathrm{out}, \mathrm{V}+2-\mathrm{m}]_\ell \cap Q_\ell^{m-1} = \emptyset$ *and* $V \leq m+1$.

Proof. We show the base cases of $m = 3, 4$. The inductive step simply repeats those arguments up to relabelling.

$m = 3$*:* Throughout, we suppose that $[\mathrm{out}, V]_\ell \subseteq Q_\ell^1$. (If not, then we are done.) From this, Lemma E.1 gives that $[\mathrm{in}]_{-\ell} \cap Q_{-\ell}^2 = \emptyset$. We divide the argument into two cases.

First, suppose that $[\mathrm{out}, \mathrm{V}-1]_{-\ell} \subseteq Q_{-\ell}^2$. In that case, any CPS strongly believes $Q_{-\ell}^2$ must assign probability one to $[\mathrm{out}, \mathrm{V}-1_{-\ell}]$ at node $V-2$. (This uses the fact that $[\mathrm{in}]_{-\ell} \cap Q_{-\ell}^2 = \emptyset$.) Thus, $[\mathrm{out}, V]_\ell$ is not a best response at node $V-2$. From this $[\mathrm{out}, V]_\ell \cap Q_\ell^3 = \emptyset$.

Second, suppose that $[\mathrm{out}, \mathrm{V}-1]_{-\ell} \cap Q_{-\ell}^2 = \emptyset$. Thus,

$$([\mathrm{out}, \mathrm{V}-1]_{-\ell} \cup [\mathrm{in}]_{-\ell}) \cap Q_{-\ell}^2 = \emptyset.$$

So, any CPS p_ℓ that strongly believes $Q_{-\ell}^2$ must have

$$p_\ell(\{s_{-\ell} : s_{-\ell}(v) = \mathrm{out}_v \text{ for some } v \leq V-3\} | S_{-\ell}) = 1.$$

(That is, *ex ante*, p_ℓ assigns probability one to the game ending at some node $v \leq V-3$, independent of the strategy that ℓ plays.) If $V \geq 5$, then there is some node $\tilde{v} \leq V-4$ at which ℓ moves and $p_\ell([\mathrm{out}, \tilde{v}+1]_{-\ell} | S_{-\ell}(\tilde{v})) = 1$. Thus, at node $\tilde{v}$, $[\mathrm{out}, \tilde{v}]_\ell$ is a unique best response. So certainly $[\mathrm{out}, V]_\ell \cap Q_\ell^3 = \emptyset$.

$m = 4$*:* Throughout, we suppose that $[\mathrm{out}, V-1]_{-\ell} \subseteq Q_{-\ell}^2$. (If not, then we are done.) From this, the base case of $m = 3$ gives that $[\mathrm{out}, V]_\ell \cap Q_\ell^3 = \emptyset$. We divide the argument into two cases.

First, suppose that $[\mathrm{out}, \mathrm{V}-2]_\ell \subseteq Q_\ell^3$. In that case, any CPS strongly believes Q_ℓ^3 must assign probability one to $[\mathrm{out}, \mathrm{V}-2]_\ell$ at node $V-3$. (This uses the fact that $([\mathrm{out}, \mathrm{V}]_\ell \cup [\mathrm{in}]_\ell) \cap Q_\ell^3 = \emptyset$.) Thus, $[\mathrm{out}, V-1]_{-\ell}$ is not a best response at node $V-3$. From this $[\mathrm{out}, V-1]_{-\ell} \cap Q_{-\ell}^4 = \emptyset$.

Second, suppose that $[\mathrm{out}, \mathrm{V}-2]_\ell \cap Q_\ell^3 = \emptyset$. Thus,

$$([\mathrm{out}, \mathrm{V}-2]_\ell \cup [\mathrm{out}, \mathrm{V}]_\ell \cup [\mathrm{in}]_\ell) \cap Q_\ell^3 = \emptyset.$$

So, any CPS $\mathrm{p}_{-\ell}$ that strongly believes Q_ℓ^3 must have

$$p_{-\ell}(\{s_\ell : s_\ell(v) = \mathrm{out}_v \text{ for some } v \leq V-4\} | S_\ell) = 1.$$

(That is, *ex ante*, $\mathrm{p}_{-\ell}$ assigns probability one to the game ending at some node $v \leq V-4$, independent of the strategy that $-\ell$ plays.) If $V \geq 6$, then there is some node $\tilde{v} \leq V-5$ at which $-\ell$ moves and $p_{-\ell}([\mathrm{out}, \tilde{v}+1]_{-\ell} | S_\ell(\tilde{v})) = 1$. Thus, at node $\tilde{v}$, $[\mathrm{out}, \tilde{v}]_{-\ell}$ is a unique best response. So certainly $[\mathrm{out}, V-2]_{-\ell} \cap Q_{-\ell}^4 = \emptyset$. ∎

Corollary E.1. *If* $V = m$*, then either* $Q_a^V = [\mathrm{out}, 1]_a$ *or* $Q_a^V = \emptyset$.

Proof. We show the result for V odd. (The case of V even is analogous.) If $Q_a^{V-2} = [\mathrm{out}, 1]_b$ or $Q_a^{V-2} = \emptyset$, then we are done. So we suppose otherwise. Observe that, by Lemmata E.1-E.2, $[\mathrm{in}]_a \cap Q_a^{V-2} = \emptyset$ and, for each $m = 3, \ldots, V-2$ odd,

$$[\mathrm{out}, V+3-m]_a \cap Q_a^{V-2} = \emptyset.$$

Thus, we must have

$$Q_a^{V-2} \in \{[\text{out}, 1]_b \cup [\text{out}, 3]_b, [\text{out}, 3]_b\}.$$

In either of these cases, $Q_b^{V-1} \in \{[\text{out}, 2]_b, \emptyset\}$. From this, it follows that $Q_a^V \in \{[\text{out}, 1]_b, \emptyset\}$. ■

Appendix F Algorithm

Proof of Proposition 8.1. Fix some game Γ. Let $\mathcal{Q} = (Q^0, Q^1, \ldots)$ be a BRP-sequence, i.e., for each finite m, $(Q^0, \ldots, Q^m)$ is an m-BRP. Since the game is finite, there is some $M(\mathcal{Q})$ so that, $Q^{M(\mathcal{Q})} = Q^{M(\mathcal{Q})+1}$. We can and do choose $M(\mathcal{Q})$ so that

$$M(\mathcal{Q}) = \begin{cases} 2\min\{|S_a|, |S_b|\} - 1 & \text{if } |S_a| \neq |S_b|, \\ 2\min\{|S_a|, |S_b|\} - 2 & \text{if } |S_a| = |S_b|. \end{cases}$$

Then take $\overline{M}$ to be the maximum of all such $M(\mathcal{Q})$ and observe that it, too, is less than or equal to $2\min\{|S_a|, |S_b|\} - 1$ (resp. $2\min\{|S_a|, |S_b|\} - 2$) if $|S_a| \neq |S_b|$ (resp. $|S_a| = |S_b|$).

It remains to show that $\overline{S}^{\overline{M}} = \overline{S}^{\infty}$. Certainly $\overline{S}^{\infty} \subseteq \overline{S}^{\overline{M}}$. Observe that that

$$\overline{S}^{\overline{M}} \subseteq \bigcup_{\text{BRP-sequences } \mathbb{S}} \overline{S}^{M(\mathcal{Q})}.$$

For each BRP-sequence $\mathcal{Q}$, $\overline{S}^{M(\mathcal{Q})} = \overline{S}^{M(\mathcal{Q})+1}$ and so $\overline{S}^{M(\mathcal{Q})}$ is itself an EFBRS. With this $\overline{S}^{M(\mathcal{Q})} \subseteq \overline{S}^{\infty}$, establishing that $\overline{S}^{\overline{M}} \subseteq \overline{S}^{\infty}$. ■

References

Battigalli, P. 1997. "On Rationalizability in Extensive Games." *Journal of Economic Theory* 74(1):40–61.

Battigalli, P. and A. Friedenberg. 2009. "Context-Dependent Forward Induction Reasoning.". http://www.public.asu.edu/~afrieden/cdfiNewWP-11-20-10.pdf.

Battigalli, P. and A. Friedenberg. 2012. "Forward Induction Reasoning Revisited." *Theoretical Economics* 7:57–98.

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

Battigalli, P. and M. Siniscalchi. 1999. "Hierarchies of Conditional Beliefs and Interactive Epistemology in Dynamic Games." *Journal of Economic Theory* 88(1):188–230.

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

Ben-Porath, E. 1997. "Rationality, Nash Equilibrium and Backwards Induction in Perfect-Information Games." *The Review of Economic Studies* 64:23–46.

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

Bradley, Stephen, Arnoldo Hax and Thomas Magnanti. 1977. "Applied mathematical programming.".

Brandenburger, A. 2003. On the Existence of a 'Complete' Possibility Structure. In *Cognitive Processes and Economic Behavior*, ed. M. Basili, N. Dimitri and I. Gilboa. Routledge pp. 30–34.

Brandenburger, A. and A. Friedenberg. 2008. "Intrinsic Correlation in Games." *Journal of Economic Theory* 141(1):28–67.

Brandenburger, A., A. Friedenberg and H.J. Keisler. 2008. "Admissibility in Games." *Econometrica* 76(2):307.

Brandenburger, Adam. 2007. "The power of paradox: some recent developments in interactive epistemology." *International Journal of Game Theory* 35(4):465–492.

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.

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

Dekel, Eddie, Marciano Siniscalchi et al. 2014. "Epistemic game theory." *Handbook of Game Theory* 4.

Friedenberg, A., W. Kets and T. Kneeland. 2016. "Bounded Reasoning: Rationality or Cognition?".

Healy, Paul J. 2015. "Preferences, Rationality, and Belief Updating in Extensive-Form Games: Experimental Evidence.".

Heifetz, A. and W. Kets. N.d. "Robust Multiplicity with a Grain of Naivite." *Theoretical Economics*. Forthcoming. forthcoming.

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.

Kuhn, H.W. 1953. "Extensive Games and the Problem of Information." *Contributions to the Theory of Games* 2(28):193–216.

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.

Reny, P. 1992. "Backward Induction, Normal Form Perfection, and Explicable Equilibria." *Econometrica* 60(3):627–649.

Rubinstein, Ariel. 2007. "Instinctive and cognitive reasoning: a study of response times." *The Economic Journal* 117(523):1243–1259.

Siniscalchi, Marciano. 2016. "Structural Rationality in Games.".

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

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.