Identification of Reasoning about Rationality *

Adam Brandenburger[†] Alexander Danieli[‡] Amanda Friedenberg[§]

April 17, 2019

Abstract

How many levels do players reason about rationality? This paper focuses on the case where the researcher observes only the strategies played. Does this information suffice? Or must the researcher look to auxiliary data relating, for example, to the players' beliefs? In simultaneous-move games, there is a well-understood method for identifying reasoning about rationality from observed play alone—one based on iterated dominance. This paper builds on the ideas from the simultaneous-move benchmark and provides a novel method for identifying reasoning about rationality in a broad class of extensive-form games. The method cannot rely on extensive-form rationalizability, an analog to the simultaneous-move benchmark of iterated dominance. This is true even for extensive-form games that are dominance solvable—that is, where extensive-form rationalizability leads to a unique solution. Instead, our method relies on a new solution concept, called an m-best response sequence. The Centipede game illustrates that there are non-trivial implications for experimental design.

1 Introduction

Strategic 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 the other player's—namely, 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, Ann 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 the other

*We thank George Mailath and three referees for many helpful comments. In addition, we have benefited from conversations with Vince Crawford, Jessy Hsieh, Terri Kneeland, Xiaomin Li, Elliot Lipnowski, and Ernesto Rivera Mora. We also thank audiences at Bristol, Oxford, Penn State, Washington University in St. Louis, McGill, National University of Singapore, Florida State University, University of Chicago, 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. A previous version of this paper was circulated under the title: "How Many Levels Do Players Reason: An Observational Challenge and Solution."

[†]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

[§]Department of Economics, University of Arizona, amanda@amandafriedenberg.org, amandafriedenberg.org

player is rational, that the other player believes that they are rational, and so on. Loosely, this benchmark involves a player reasoning, at all levels, about the other player's rationality. But, in practice, players may depart from such a benchmark.

This raises the question: Can levels of reasoning about rationality be inferred from behavior alone? Toward that end, we focus on the case where the researcher only observes the strategies played. Does this information suffice? Or must the researcher look to auxiliary data—e.g., data about the players' beliefs?

In simultaneous-move games, there is a well-understood method for identifying reasoning about rationality from behavior alone. That method is based on rationalizability (or, equivalently, iterated dominance). Dominance-solvable games are well-suited to implement the method. Below, we review why.

This paper builds on the ideas from the simultaneous-move benchmark and provides a novel method for identifying reasoning about rationality in a broad class of extensive-form games. Unlike the simultaneous-move benchmark, the method cannot rely on extensive-form rationalizability. As a consequence, the researcher may not be able to infer levels of reasoning about rationality even in a game that is dominance solvable—that is, in a game where extensive-form rationalizability leads to a unique prediction. We will draw out implications for the interpretation of several experimental results. Let us preview the ideas.

Simultaneous-Move Benchmark In simultaneous-move games, there is an accepted way of formalizing "reasoning about rationality" and a well-understood method for identifying such reasoning from observed behavior. We begin with that benchmark.

In the benchmark case, reasoning about rationality is 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. She *believes* the other player is rational if her belief assigns probability 1 to the other player's rationality. She is *rational and* 1^{st}*-order believes rationality* if she is rational and believes the other player is rational. Inductively, she is rational and m^{th}-order believes rationality if she is rational and n^{th}-order believes rationality, for all $n = 1, \ldots, m-1$.

The concept of rationality and m^{th}-order belief of rationality is formalized within an epistemic framework. (The framework is suitable to do so, since it describes the players' hierarchies of beliefs about how the game is played. See, e.g., Brandenburger, 2007 and Dekel and Siniscalchi, 2014 for an overview.) The $(m+1)$*-rationalizable* strategies (Bernheim, 1984; Pearce, 1984) capture the behavior associated with rationality and m^{th}-order belief of rationality. 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 Brandenburger and Friedenberg, 2008 for a proof.) This standard result can be used to provide a *rationality bound*: 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, $(m-1)^{th}$-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.

In light of this result, dominance-solvable games are well-suited to infer levels of reasoning about rationality. In such games, each round of rationalizability eliminates at least one strategy for some player, until no further eliminations are possible. At that point, there is a unique prediction. Thus, at most one strategy of each player is consistent with rationality and common belief of rationality.

To sum up, in simultaneous-move games, we can use behavior alone to infer a rationality bound. This can be achieved by inferring the maximum m so that the behavior is m-rationalizable.

Beyond Simultaneous-Move Games Let's highlight the architecture of the simultaneous-move game analysis: First, we specify an epistemic framework. Within that framework, we formulate what it means to reason about rationality; this is formalized as rationality and m^{th}-order belief of rationality. Finally, as a theorem, we derive behavioral implications of rationality and m^{th}-order belief of rationality. This came in the form of the $(m+1)$-rationalizable strategies.

One might wonder if it is really necessary to take a formal epistemic approach. In simultaneous-move games, one might intuit 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 analysis. Perhaps, more generally, the epistemic approach is not needed? There are two natural conjectures. First, in the context of perfect-information games, we can skip the step of going through an epistemic approach by invoking the backward-induction algorithm. Second, in more general extensive-form games, we can instead invoke extensive-form rationalizability (EFR, Pearce 1984). However, neither concept can be used to provide a bound on reasoning about rationality—at least not in an obvious way. Here, we explain why the backward-induction algorithm cannot be used. Examples 3.2-3.3 will explain why EFR cannot be used.

Consider the game in Figure 1.1. Suppose we observe a subject in the role of Ann play I_a-o_a. We should conclude that this subject is irrational: The strategy I_a-o_a is dominated by O_a. As such, there is no belief that Ann can hold that would make I_a-o_a a best response for Ann.

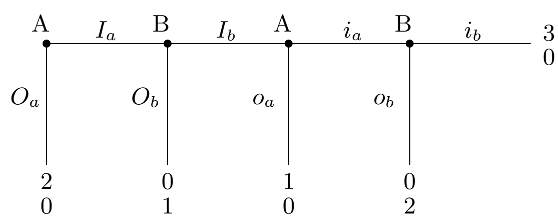

Figure 1.1: Reny (1992)

That said, the backward-induction algorithm cannot easily be used to conclude that the subject is irrational. No strategy of Ann 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. Thus, natural attempts to use the backward-induction algorithm to infer a rationality bound would be misleading.

The fact that we cannot use the backward-induction algorithm to obtain a bound on the level of reasoning about rationality arises from a fundamental feature of backward induction. On the first round of the procedure, the algorithm need not eliminate all (conditionally) dominated strategies, i.e., all strategies that are not a sequential best response. Thus, our conclusion does not rely on details of epistemic assumptions—that is, it is independent of exactly how "reasoning about rationality" is formalized.

This Paper We have argued for the need to follow the structure of the simultaneous-move benchmark. Therefore, we begin by specifying an epistemic framework appropriate for studying extensive-form games. Within this framework, we formalize what it means to "reason about rationality." We take this to mean *rationality and m^{th}-order strong belief of rationality* (Battigalli and Siniscalchi, 2002). (The choice is discussed in Sections 3-9.c.) We provide a behavioral characterization of rationality and m^{th}-order strong belief of rationality in a particular class of games: *generic games* (Definition 6.4). The characterization rests on a new solution concept called an *m-best response sequence* (m-BRS, Definition 6.3). Specifically:

> **Theorem 6.1:** Fix a generic game. A strategy is consistent with rationality and m^{th}-order strong belief of rationality if and only if the strategy is consistent with an m-BRS.

The strategies that survive m rounds of EFR form an m-BRS. But the converse does not hold. Thus, attempts to use EFR to infer levels of reasoning about rationality will be misguided. In fact, there are many games which are dominance solvable (according to EFR), but which only permit what we will call *trivial identification*: the ability to separate behavior that is irrational from behavior that is consistent with rationality and common strong belief of rationality. These games are not well-suited for inferring rationality bounds, since there are no strategies that are consistent with rationality and m^{th}-order—but not $(m+1)^{th}$-order—strong belief of rationality. This is an important point, since the literature has pointed to behavior in such games as providing evidence on the limits of strategic reasoning. (See, e.g., Cooper, DeJong, Forsythe and Ross 1993, Brandts and Holt 1995, Huck and Müller 2005, and Brandts, Cabrales and Charness 2007.)

Games that permit trivial identification pose a problem for the researcher who attempts to infer reasoning about rationality from a dataset that consists only of information about the strategies played. If *all* games permit trivial identification, the researcher would not be able to use information about the strategies played to infer reasoning about rationality (i.e., in *any* game).

This raises the question: Are there extensive-form games which permit *non-trivial identification*? We will see that the Centipede game is one such example. The researcher can use behavior of the first mover to disentangle rationality and m^{th}-order strong belief of rationality from rationality and $(m+2)^{th}$-order strong belief of rationality. However, there is an important limitation. The behavior of the second mover only permits trivial identification. Thus, the researcher can use first-mover—but not second-mover—data for the purposes of identification.

Can the researcher do better? Can the researcher design an experiment—on a game with extensive structure—that permits tighter identification? Perhaps the researcher can design an experiment that permits non-trivial identification for all players. Or, perhaps there is an experiment that can separate rationality and m^{th}-order strong belief of rationality from rationality and $(m+1)^{th}$-order strong belief of rationality. We leave the potential for such tighter identification as an open question. (Though, see Section 9.g for a preliminary discussion.)

Importantly, we don't see this paper as providing the 'ultimate' experimental design. Instead, we hope that the results will be a useful step to bridge the gap between epistemic game theory and experimental economics. In fact, the paper fits in a growing literature on bridging the gap in the context of extensive-form games. For instance, Healy (2017) seeks to elicit conditional beliefs about rationality. Siniscalchi (2016) seeks to elicit conditional hierarchies of beliefs about the strategies played. By contrast, here, we focus on identification of the rationality bound when the researcher has access to data only about behavior.

Reasoning about Rationality vs. Other Limitations This paper focuses on *bounded reasoning about rationality*. It is important to distinguish this idea from two other limitations on "reasoning" discussed in the literature: *limited foresight* and *limited ability to engage in interactive reasoning*.

Limited foresight is the idea that players may be unable to look more than k steps forward into the tree, to understand what players may or may not choose in the future. This is an important bound on how players reason—one which is not yet fully understood. (See Mantovani, 2014 for an early attempt.) This paper abstracts away from limits on foresight. Arguably, many analyses of experimental data implicitly do the same. (One may think of this as assuming the researcher has designed an experiment that is sufficiently simple so that limited foresight is not a significant issue.)

Limited ability to engage in interactive reasoning reflects limits on the ability to compose sentences of the form "I think that you think that I think ...."[1] There is evidence from cognitive psychology that subjects are limited in their ability to engage in interactive reasoning. (See, e.g., Kinderman, Dunbar and Bentall, 1998 and Stiller and Dunbar, 2007, among others.) Note, carefully, limited ability to engage in interactive reasoning can exist independent of reasoning about rationality. (For instance, Ann may find it difficult to produce the sentence "Bob believes that I believe the sun rises at 6am.") But, limited ability to engage in interactive reasoning can limit reasoning about rationality.

We take no position on *why* a player may exhibit bounded reasoning about rationality. It may well arise from a limited ability to engage in interactive reasoning. But, it need not. 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 may be irrational. Or, by reasoning that "Bob is rational but believes she might be irrational." Etc. Section 9.b discusses the significance of the source of the rationality bound.

Reasoning about Rationality vs. Level-k Reasoning There is a long tradition of using behavior to infer levels of iterative reasoning. The literature typically organizes the data by appealing to one of two prominent models: the level-k model or the cognitive hierarchy model. (See Nagel 1995, Stahl and Wilson 1995, Costa-Gomes, Crawford and Broseta 2001, Camerer, Ho and Chong 2004, and Costa-Gomes and Crawford 2006, among others.) Many applications of these models involve some amount of *reasoning about irrationality*. To understand why, note that, in these models, level-0 types are associated with a mixed strategy and that mixed strategy pins down the beliefs of level-1 types. Often, the level-0 mixed strategy is chosen to assign positive probability to a dominated—and, therefore, irrational—strategy. When that occurs, level-1 types have a belief that assigns positive probability to irrationality.[2] Thus, the level-k and cognitive hierarchy models are conceptually distinct from iterative reasoning about rationality.

That said, in simultaneous-move games, there is an important connection in terms of predicted behavior. There, level-k behavior is consistent with k-rounds of rationalizability.[3] So, despite the fact that level-k *reasoning* differs from rationality and $(k-1)^{th}$-order belief of rationality, it generates *behavior* that is consistent with rationality and $(k-1)^{th}$-order belief of rationality. As a consequence, a subject who would

[1] Camerer, Ho and Chong (2004) appear to use the phrase "reasoning" (unqualified) to capture this idea. Some other papers have followed suit; however, this choice is not uniform across the literature. For instance, Costa-Gomes, Crawford and Broseta (2001) use the phrase to capture limited rounds of iterated dominance. Since the phrase "reasoning" (unqualified) can reflect many ideas, we prefer to be specific and, so, only use the phrase with qualification.

[2] This is not always the case. A prominent example is Arad and Rubinstein (2012), where the level-0 behavior is not dominated. (In fact, it is rationalizable.) So, there, level-1 types can be viewed as types that believe rationality.

[3] The same does not hold for the cognitive hierarchy model: A 2-cognitive hierarchy type need not be 2-rationalizable. (On this point, Heifetz and Kets's (2018) notion of rationalizability is behaviorally similar to the cognitive hierarchy model.)

be identified as a level-k reasoner (for $k \geq 1$) will—in the terminology of this paper—have an identified rationality bound of $m \geq k$. The fact that m may be strictly larger than k reflects the fact that we identify the rationality bound as the maximum m so that rationality and m^{th}-order belief of rationality is consistent with observed behavior. The fact that the level-k model is able to pin down an exact level—instead of a maximum—reflects the fact that the model imposes auxiliary assumptions on the players' hierarchies of beliefs: It imposes an assumption about level-0 behavior and that pins down level-1's first-order beliefs about the strategies played.[4] This pins down level-2's second-order beliefs. Etc.

This highlights an important point: Unlike the level-k and cognitive hierarchy literatures, we do not impose auxiliary assumptions on players' hierarchies of beliefs (or, even, the structures within which they lie). In this respect, our approach is closer in spirit to Kneeland (2015). In the context of a specific experimental design, Kneeland identifies rationality and m^{th}-order belief of rationality from observed behavior—without imposing auxiliary assumptions on players' hierarchies of beliefs.[5]

Why not impose auxiliary assumptions on hierarchies of beliefs? Examples 3.2-3.3 will highlight the fact that, in the context of extensive-form games, imposing such assumptions can have substantial consequences—consequences that do not arise in the context of simultaneous-move games. As such, we seek non-trivial identification absent such auxiliary assumptions.

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 want to identify the maximum level reasoning about rationality consistent with the observed behavior. The starting point is to formally define what we mean by "reasoning about rationality." This will involve describing the strategic situation by, what is called, an epistemic game. An epistemic game will consist of the game itself (here, BoS with an Outside Option) and an epistemic type structure.

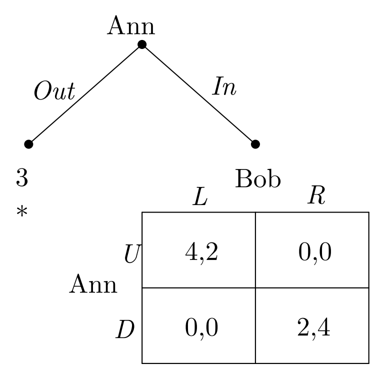

Figure 2.1: Battle of the Sexes with an Outside Option

Why include an epistemic type structure in the description of the strategic situation? Observe that the strategy *Out* is rational—i.e., a best response—if Ann believes that Bob will play R; it is irrational if

[4] 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.

[5] Kneeland's exclusion restriction is an assumption about the process by which beliefs are formed.

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 need to specify both Ann's belief about the strategy Bob plays and 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. An epistemic type structure is a framework that implicitly describes these hierarchies of beliefs.

Extensive-Form Game Write Γ for a finite two-player extensive-form game (Kuhn, 1953) with perfect recall and without moves by nature. The players of the game are a (Ann) and b (Bob).[6] 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 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 c 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 $s_c \in S_c(h)$, 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 $\text{proj}_{S_a} Q = \text{proj}_{S_b} Q = \emptyset$.

Type Structure The underlying premise is that players face uncertainty about the play of the game. Thus, players will have beliefs about the strategies the others play. 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, he 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 need to specify conditional beliefs about the play of the game. We will use an epistemic type structure to model the players' hierarchies of conditional beliefs about the play of the game.

Refer to $(\Omega, \mathcal{S}(\Omega))$ as a probability space, when Ω is a compact metric space and $\mathcal{S}(\Omega)$ is the Borel σ-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 $(\Omega, \mathcal{S}(\Omega), \mathcal{E})$ a **conditional probability space** if $(\Omega, \mathcal{S}(\Omega))$ is a probability space and $\mathcal{E} \subseteq \mathcal{S}(\Omega)\backslash\{\emptyset\}$ is finite. The collection $\mathcal{E}$ is a set of **conditioning events**. Since $\mathcal{S}(\Omega)$ is clear from the context, we suppress reference to $\mathcal{S}(\Omega)$ 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 $\mathrm{p} = (p(\cdot|E) : E \in \mathcal{E})$ so that, for each $E \in \mathcal{E}$, $p(\cdot|E) \in \mathcal{P}(\Omega)$ with $p(E|E) = 1$.

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

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

An array p specifies a belief, viz. $p(\cdot|E)$, for each conditioning event E. We refer to the beliefs $p(\cdot|E)$ as conditional beliefs. If the array is a CPS, then the conditional beliefs must satisfy the rules of conditional probability when possible. 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. $S_b(\phi) = S_b$. 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 is 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 $(S_a, \mathcal{E}_b)$. Etc.

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 example informally illustrates such an assumption.

Example 2.1. Consider BoS with an Outside Option. Suppose it is 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. 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.

We can capture this restriction on beliefs by a type structure $\mathcal{T} = (\Gamma; T_a, T_b; \beta_a, \beta_b)$ that satisfies the following properties:

- Each $\beta_a(t_a)(\cdot|S_b \times T_b)$ assigns probability 1 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 models the assumption that there is common (full) belief that Bob plays $\{R\}$. (See Example 3.1 on full belief.) □

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., $\beta_c : T_c \to \mathcal{P}(S_{-c} \times T_{-c})$. □

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 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.

3 Epistemic Conditions

The set of states associated with a type structure $\mathcal{T}$ describes the set of possible strategy-type pairs that can obtain. This imposes no restrictions on behavior or how players reason about rationality. The epistemic conditions will do that. There are two types of epistemic conditions: First, we impose the behavioral condition of rationality. Second, we restrict beliefs so that they arise from reasoning about rationality.

Rationality Fix some $X_c \subseteq S_c$. Say $s_c \in X_c$ is **optimal under** $\mu_c \in \mathcal{P}(S_{-c})$ **given** X_c if

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

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 optimal under $p_c(\cdot | S_{-c}(h))$ given $S_c(h)$.

Write $\mathbb{BR}[\mathrm{p}_c]$ for the set of strategies s_c that are a sequential best response under $\mathrm{p}_c \in \mathcal{A}(S_{-c}, \mathcal{E}_c)$. So, each $s_c \in \mathbb{BR}[\mathrm{p}_c]$ is optimal under each of the conditional beliefs $p_c(\cdot | S_{-c}(h))$, provided h is an information set allowed by s_c.

Each $\beta_c(t_c)$ induces a CPS in $\mathcal{C}(S_{-c}, \mathcal{E}_c)$ via marginalization. The marginal CPS, viz. $\mathrm{marg}_{S_{-c}} \beta_c(t_{-c})$, is a CPS $\mathrm{p}_c \in \mathcal{C}(S_{-c}, \mathcal{E}_c)$ with $p_c(\cdot | S_{-c}(h)) = \mathrm{marg}_{S_{-c}} \beta_c(t_c)(\cdot | S_{-c}(h) \times T_{-c})$ for each $S_{-c}(h) \in \mathcal{E}_c$.

Definition 3.2. Say (s_c, t_c) is **rational** if $s_c \in \mathbb{BR}[\mathrm{marg}_{S_{-c}} \beta_c(t_c)]$.

So, (s_c, t_c) is rational if s_c is a sequential best response under the marginal CPS $\mathrm{marg}_{S_{-c}} \beta_c(t_c)$.

Reasoning about Rationality We next impose the requirement that a player reasons about the other player's rationality. We take "reasoning about rationality" 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}.

Strong belief asks that a type maintain a hypothesis so long as it is not contradicted by observed play. Thus, it requires that a type rationalize past behavior when possible. In this sense, it captures forward induction reasoning. (See Battigalli and Siniscalchi 2002 and Battigalli and Friedenberg 2012 for a more complete discussion.)

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 1 to the event E_{-c}. To better understand 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.[7]

Example 3.1. Consider the game in Figure 3.1. Observe that, in this game, Ann's strategy 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 $s_b = O_b$.

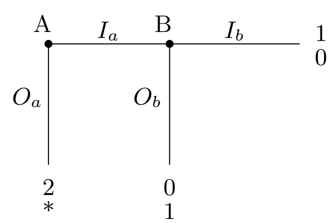

Figure 3.1: Strong Belief versus Full Belief

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 1 to E_a given every conditioning event. For any type structure, there is no type of Bob that fully believes that Ann is rational. This is because, conditional upon Bob reaching the information set associated with Ann's choice of I_a, Bob must assign probability 1 to the event that Ann is irrational.

Suppose we were to take "reasoning about rationality" to be 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 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 choose an irrational move. It does not reflect a lack of reasoning about rationality on the part of Bob.

In this scenario, strong belief of rationality would only require that a type of Bob initially assign probability 1 to Ann's rationality; it would not require that Bob maintain a hypothesis that Ann is rational, when he observes evidence inconsistent with Ann's rationality. In fact, we can construct a type structure and a state at which there is "rationality and common strong belief of rationality." (The concept will be defined shortly.) □

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

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 $R_b^m(\mathcal{T})$ by

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

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 $R^\infty(\mathcal{T}) = \cap_{m \geq 1} R^m(\mathcal{T})$ 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 $\mathcal{T}$. A natural conjecture is that this gives rise to the set of extensive-form rationalizable (EFR, Pearce, 1984) strategies. 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 EFR. However, this need not be the case for a type-incomplete structure $\mathcal{T}$. In that case, the predictions of RmSBR may be disjoint from the 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 sequential best response—i.e., a CPS that assigns probability 1 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 1 to R (resp. L). But there is no CPS so that *In*-D is a sequential best response. Thus, one round of EFR gives

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

Now observe that a CPS on S_a that strongly believes $\text{EFR}_a^1 = \{Out, In\text{-}U\}$ must assign probability 1 to *In*-U conditional on BoS; L is the unique sequential best response under any such CPS. As such,

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

With this, *In*-U is the unique sequential best response under any CPS on S_b that strongly believes $\text{EFR}_b^2 = \{L\}$. Thus,

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

So, 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,

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

So, the EFR predictions are also RmSBR predictions, provided the type structure 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 understood that "Bob is a bully." Now

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

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 1 to $\{Out\} \times T_a$ and, conditional upon Ann playing *In*, assigns probability 1 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 1 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 Identified Rationality Bound

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 deduce the subject's actual rationality bound. For instance, suppose the type structure is $\mathcal{T}$ and the true state is

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

Then Ann's *actual rationality bound* is 2: She is rational, strongly believes Bob is rational, but does not strongly believe "Bob is rational and strongly believes I am rational." If the researcher knew $\mathcal{T}$ and the state, the researcher would know Ann's actual rationality bound.

Here, we focus on the case where the researcher cannot observe the type structure and the players' hierarchies of beliefs (i.e., types). That is, we suppose that the researcher *only* observes the game Γ and the strategies played. As such, we focus on the maximum level of RmSBR consistent with observed behavior. Because the researcher does not observe the players' hierarchies of beliefs (i.e., types) and the type structure within which they lie, this maximum is taken over all type structures and beliefs.

To understand what this involves, return to BoS with an Outside Option. Suppose the researcher observes Ann play *Out*. If the researcher knew the players' type structure but not Ann's actual hierarchy of beliefs, he could use the type structure to provide a bound on the level of RmSBR consistent with *Out*. For instance, if the researcher knew that the players' type structure were type-complete, then the researcher would infer that Ann's behavior is consistent with *at most* R1SBR. However, if the researcher knew that it was understood that "Bob is a bully," the researcher would instead infer that Ann's behavior is consistent with RCSBR. Because the researcher does not know the players' type structure, we take the maximum across both type structures and beliefs. So, if we observe Ann play *Out*, we will identify her behavior as consistent with RCSBR. We now make this idea precise.

Identified Rationality Bound We think of the researcher as having access to a dataset that reflects the strategies played in the role of player c; write $\mathbb{D}_c \subseteq S_c$ for this dataset.[8] For each finite $m \geq 0$, say the data $s_c \in \mathbb{D}_c$ is **consistent with RmSBR** if there exists some type structure $\mathcal{T}$ with $s_c \in \text{proj}_{S_c} R_c^{m+1}(\mathcal{T})$.

[8] Section 9.f discusses the case of a sparser dataset—e.g., where the researcher only observes a signal of the strategies played. Section 9.g discusses the case of a richer dataset—where the researcher has observations across games.

Likewise, say the data $s_c \in \mathbb{D}_c$ is **consistent with RCSBR** if there exists some type structure $\mathcal{T}$ with $s_c \in \text{proj}_{S_c} R_c^\infty(\mathcal{T})$.

Definition 4.1.

(1) Say $s_c \in \mathbb{D}_c$ has an **identified rationality bound of** 0 if s_c is not consistent with R0SBR.

(2) For finite $m \geq 1$, say $s_c \in \mathbb{D}_c$ has an **identified rationality bound of** m if s_c is consistent with R$(m-1)$SBR but not consistent with RmSBR.

(3) Say $s_c \in \mathbb{D}_c$ has an **identified rationality bound of** ∞ if, for each finite $m \geq 1$, s_c is consistent with RmSBR.

Suppose the researcher observes data $s_c \in \mathbb{D}_c$. The identified rationality bound is $m < \infty$ if (i) there is a type structure $\mathcal{T}^m$ with $s_c \in \text{proj}_{S_c} R_c^m(\mathcal{T}^m)$, and (ii) for each type structure $\mathcal{T}$, $s_c \notin \text{proj}_{S_c} R_c^{m+1}(\mathcal{T})$. The identified rationality bound is ∞ if, for each finite $m \geq 1$, there exists a type structure $\mathcal{T}^m$, with $s_c \in \text{proj}_{S_c} R_c^m(\mathcal{T}^m)$. (Anticipating the discussion in Section 9.b, more precise terminology would be an *identified rationality bound of* m *for* Γ. We drop reference to Γ for verbal simplicity.) There are two important subtleties that deserve discussion.

Remark 4.1. Suppose we observe a subject play s_c and the *identified* rationality bound for s_c is m. The subject's *actual* rationality bound cannot be strictly greater than m. But, it may, in fact, be strictly less than m. That is, the identified rationality bound is an upper bound on the subject's actual rationality bound.

As an illustration, refer back to BoS with an Outside Option. By Example 3.3, there exists some type structure $\mathcal{T}$ and some type t_a, so that $(s_a, t_a) \in \bigcap_m R_a^m(\mathcal{T})$. Thus, if the researcher observes the subject play *Out*, the identified rationality bound is ∞. But, suppose, the subjects' actual type structure is type-complete. Then, by the analysis in Example 3.2, the subject's actual rationality bound must be $m \leq 2$. In fact, it may well be the case that the subject's rationality bound is $m < 2$.

To see this last claim, note that in a type-complete structure, there exists types t_a and t_b that satisfy the following criteria:

$$\beta_b(t_b)(\{\textit{In-U}\} \times T_a \mid S_a \times T_a) = 1 \qquad \text{and} \qquad \beta_a(t_a)((R, t_b) \mid S_a \times T_a) = 1.$$

Note, (R, t_b) is irrational. (Conditional upon t_b's information set being reached, he must continue to believe that Ann will play U; with this, L is the unique sequential best response for t_b.) So, at the start of the game, t_a does not believe that Bob is rational. As such, $(\textit{Out}, t_a)$ is rational, but t_a does not strongly believe that Bob is rational. Put together, if a subject plays *Out*, but her actual beliefs conform to those of t_a, then her actual rationality bound is 1. □

Remark 4.2. If the data is consistent with RCSBR, then the data has an identified rationality bound of ∞. 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 RCSBR (i.e., there exists some type structure $\mathcal{T}$ and some $(s_c, t_c) \in R_c^\infty(\mathcal{T})$).

Despite the above, in Section 9.e, we will show that the converse also holds: If the data has an identified rationality bound of ∞, then the data is consistent with RCSBR. More informally, if there is no identified finite rationality bound, then the data must, in fact, be consistent with RCSBR. □

We seek to construct a partition

$$\mathcal{IB}_c = \{\mathrm{IB}_c^0, \mathrm{IB}_c^1, \ldots, \mathrm{IB}_c^m, \ldots, \mathrm{IB}_c^\infty\}$$

on the strategies of player c, viz. S_c, so that the following holds: For each $m = 0, 1, 2, \ldots, \infty$, $s_c \in \mathrm{IB}_c^m$ if and only if s_c would be identified as having a rationality bound of m were it observed in the data. We next discuss a challenge in doing so.

A Challenge There is a natural approach to constructing the partition $\mathcal{IB}_c$. Refer to Figure 4.1. For each Γ-based type structure $\mathcal{T}$, define $S_c^m(\mathcal{T}) := \mathrm{proj}_{S_c} R_c^m(\mathcal{T})$, i.e., the projection of the strategy-type pairs in $R_c^m(\mathcal{T})$ onto S_c. Then set

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

Note, $\overline{S}_c^m$ is the collection of strategies s_c consistent with R$(m-1)$SBR (in some type structure). So, for each finite $m \geq 0$, $\mathrm{IB}_c^m = \overline{S}_c^m \backslash \overline{S}_c^{m+1}$. Moreover, $\mathrm{IB}_c^\infty = S_c \backslash \bigcup_{m \geq 0} \mathrm{IB}_c^m$.

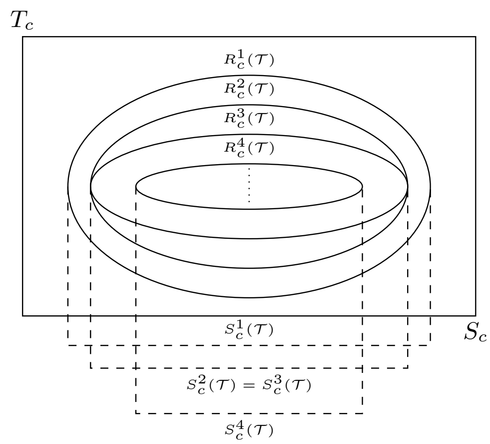

Figure 4.1: Projections of $R_c^m(\mathcal{T})$

In practice, there is a challenge in implementing this approach: We defined the set $\overline{S}_c^m$ as the union of the sets $S_c^m(\mathcal{T})$ over all type structures $\mathcal{T}$ associated with Γ. So, this step requires searching across all type structures $\mathcal{T} \in \mathbb{T}(\Gamma)$, which is an uncountable set.

To overcome the issue, we will seek to identify $\overline{S}_c^m$ from properties of the game Γ alone. Toward that end, fix a finite m and a subset of strategies $Q_a \times Q_b \subseteq S_a \times S_b$. We would like a test—based only on the game itself—so that $Q_a \times Q_b$ passes the test if and only if there exists a type structure $\mathcal{T} \in \mathbb{T}(\Gamma)$ so that $Q_a \times Q_b = S^m(\mathcal{T}) := \mathrm{proj}_{S_a \times S_b} R^m(\mathcal{T})$: In words, $Q_a \times Q_b$ passes the test if and only if it is the set of R$(m-1)$SBR predictions for some type structure $\mathcal{T}$.

An analogous test is known for RCSBR. Specifically, $Q_a \times Q_b$ is an extensive-form best response set (EFBRS) if and only if there exists a type structure $\mathcal{T}$ so that $Q_a \times Q_b$ is the set of RCSBR predictions for $\mathcal{T}$. (See Battigalli and Friedenberg, 2012.) So, $Q_a \times Q_b$ will pass the RCSBR-test if it is an EFBRS and fail the RCSBR-test otherwise. With this in mind, we will use the solution concept of EFBRS as a benchmark, by which we think of addressing the broader issue. We begin by reviewing the concept.

5 The EFBRS Benchmark

Fix a type structure and $\mathcal{T}$ and let $S^\infty(\mathcal{T}) := \text{proj}_{S_a \times S_b} R^\infty(\mathcal{T})$ be the set of RCSBR predictions for $\mathcal{T}$. Note, $S^\infty(\mathcal{T}) = S_a^\infty(\mathcal{T}) \times S_b^\infty(\mathcal{T})$ is a product set. Fix a predicted strategy for Ann, viz. $s_a \in S_a^\infty(\mathcal{T})$, and a type t_a so that $(s_a, t_a) \in R_a^\infty(\mathcal{T})$. Write p_a for the marginal CPS of t_a, i.e. $\text{p}_a = \text{marg}_{S_b}\beta_a(t_a)$. Observe that Ann's strategy and the marginal CPS, viz. (s_a, p_a), must satisfy three properties: First, s_a must be a sequential best response under p_a. (This follows from the fact that (s_a, t_a) is rational.) Second, p_a must strongly believe Bob's RCSBR prediction for $\mathcal{T}$, viz. $S_b^\infty(\mathcal{T})$. (This follows from the fact that t_a must strongly believe $R_b^\infty(\mathcal{T})$.) Finally, if r_a is also a sequential best response under p_a, r_a must be contained in $S_a^\infty(T)$, i.e., r_a must be one of Ann's RCSBR predictions for $\mathcal{T}$. (This follows from the following property: If (r_a, t_a) is rational and (s_a, t_a) satisfies RCSBR for $\mathcal{T}$, then (r_a, t_a) satisfies RCSBR for $\mathcal{T}$.) This last property can be viewed as a maximality condition.

These three properties motivate the definition of an EFBRS.

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 $\text{p}_c \in \mathcal{C}(S_{-c}, \mathcal{E}_c)$ so that the following hold:

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

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

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

EFBRS's are defined by properties on the game itself. In particular, an EFBRS is a subset of strategies, viz. $Q_a \times Q_b$, that satisfies a certain fixed point requirement: For each $s_a \in Q_a$, there exists a CPS p_a defined only on the strategies of Bob so that (i) s_a is a sequential best response under p_a, (ii) p_a strongly believes Bob's prediction Q_b, and (iii) Q_a satisfies a requisite maximality property. These correspond to the properties derived from RCSBR. This reflects the fact that the EFBRS concept characterizes the RCSBR predictions across all type structures.

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* $\text{proj}_S R^\infty(\mathcal{T}) = Q_a \times Q_b$.

As a corollary:

Corollary 5.1. *For each game* Γ,

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

Corollary 5.1 allows us to identify the strategies consistent with RCSBR based on properties of the game alone. We now illustrate how this is implemented in BoS with an Outside Option.

Example 5.1. Return to Figure 2.1. We will show that

$$\overline{S}^\infty = \Big(\{Out\} \times \{L, R\}\Big) \cup \Big(\{In\text{-}U\} \times \{L\}\Big).$$

First, each of $\{Out\}\times\{R\}$, $\{Out\}\times\{L,R\}$, and $\{In\text{-}U\}\times\{L\}$ are EFBRS's. Take $Q_a\times Q_b=\{Out\}\times\{R\}$: *Out* is a unique sequential best response under a CPS that, at each information set, assigns probability 1 to $Q_b=\{R\}$; R is a unique sequential best response under a CPS that initially assigns probability 1 to $Q_a=\{Out\}$ and, conditional upon BoS, assigns probability 1 to $\{In\text{-}D\}$. And so on.

In fact, these are the only three non-empty EFBRS's. To see this, observe that, at the start of the game, *In-D* is dominated; by Property (1), it cannot be contained in any EFBRS. So, if $Q_a\times Q_b$ is an EFBRS with $In\text{-}U\in Q_a$, then $Q_b=\{L\}$. (This follows from Properties (1)-(2).) Moreover, if $Q_a\times Q_b=Q_a\times\{L\}$ is an EFBRS, then $Out\notin Q_a$. (Again, use Properties (1)-(2).) With these facts, there are no additional non-empty EFBRSs. Thus, the assertion follows from Corollary 5.1. □

6 The m-BRS

The EFBRS concept can be viewed as a collection of sets, each of which satisfy a certain fixed point property: If s_a is contained in Ann's solution Q_a, then s_a is a sequential best response under a CPS that strongly believes Bob's solution Q_b. The EFBRS inherits this fixed point property from RCSBR itself. If $(s_a,t_a)\in R_a^\infty(\mathcal{T})$, then t_a strongly believes an event of the same order, namely $R_b^\infty(\mathcal{T})$.

To obtain a finite-order analogue, we need to depart from this fixed point property—converting it into an iterative property. This is because RmSBR is not a fixed point concept: If $(s_a,t_a)\in R_a^3(\mathcal{T})\backslash R_a^4(\mathcal{T})$, then t_a does not strongly believe the event of the same order $R_b^3(\mathcal{T})$. Instead, t_a strongly believes the lower-order events $R_b^0(\mathcal{T}),R_b^1(\mathcal{T}),R_b^2(\mathcal{T})$. More generally, if $(s_a,t_a)\in R_a^m(\mathcal{T})\backslash R_a^{m+1}(\mathcal{T})$, then t_a strongly believes the lower-order events $R_b^0(\mathcal{T}),\ldots,R_b^{m-1}(\mathcal{T})$. We will build off this fact to convert the EFBRS concept to an iterative property. That property will apply to a decreasing sequence of product sets.

Definition 6.1. Say $(Q^0,\ldots,Q^m)$ is a **decreasing sequence of product sets** 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 6.2. 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 product sets 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 6.2 appears similar to Definition 5.1. The central difference arises in condition (BRP.2). Instead of the CPS strongly believing X_{-c}, the CPS strongly believes the lower-order sets $Q_{-c}^0,\ldots,Q_{-c}^m$. (Note, $X_{-c}\subseteq Q_{-c}^m\subseteq Q_{-c}^{m-1}\subseteq\cdots\subseteq Q_{-c}^0$.)

Definition 6.3. Let $m\geq 1$. Say $(Q^0,\ldots,Q^m)$ forms an m-**(extensive-form) best response sequence** (m-**BRS**) if $Q^1\neq\emptyset$ and, for each $n=1,\ldots,m-1$, Q^{n+1} satisfies the best response property relative to $(Q^0,\ldots,Q^n)$.

Remark 6.1. For each $m\geq 2$, $(Q^0,\ldots,Q^m)$ is an m-BRS if and only if (i) $(Q^0,\ldots,Q^{m-1})$ is an $(m-1)$-BRS and (ii) Q^m satisfies the best response property relative to $(Q^0,\ldots,Q^{m-1})$. □

A 1-BRS is some $(Q^0, Q^1) = (S_a \times S_b, Q_a^1 \times Q_b^1)$, where

$$Q_c^1 = \bigcup_{\mathrm{p}_c \in E_c} \mathbb{BR}[\mathrm{p}_c]$$

for some non-empty $E_c \subseteq \mathcal{C}(S_{-c}, \mathcal{E}_c)$. An $(m+1)$-BRS is some $(Q^0, \ldots, Q^m, Q^{m+1})$, where $(Q^0, \ldots, Q^m)$ is an m-BRS and Q^{m+1} satisfies the best response property relative to $(Q^0, \ldots, Q^m)$. Thus, it is an iterative procedure that is a natural analogue of the EFBRS. In fact:

Proposition 6.1. *For each $\mathcal{T}$, the sequence $(\mathrm{proj}_S R^0(\mathcal{T}), \ldots, \mathrm{proj}_S R^m(\mathcal{T}))$ forms an m-BRS.*

Say Q is **consistent with the m-BRS** if there exists some $(m-1)$-BRS, viz. $(Q^0, \ldots, Q^{m-1})$, so that Q satisfies the extensive-form best response property relative to $(Q^0, \ldots, Q^{m-1})$. By Proposition 6.1,

$$\overline{S}^m := \bigcup_{\mathcal{T} \in \mathbb{T}(\Gamma)} \mathrm{proj}_S R^m(\mathcal{T}) \subseteq \bigcup_{Q \text{ is consistent with the m-BRS}} Q.$$

That is, the union over the m-BRS's provide an upper bound on the behavior consistent with R$(m-1)$SBR across all type structures.[9]

A natural analogue to Corollary 5.1 is that

$$\overline{S}^m := \bigcup_{\mathcal{T} \in \mathbb{T}(\Gamma)} \mathrm{proj}_S R^m(\mathcal{T}) = \bigcup_{Q \text{ is consistent with the m-BRS}} Q, \tag{1}$$

i.e., $\overline{S}^m$ is the union over sets Q that are consistent with the m-BRS. In fact, there is a natural conjecture that would imply Equation (1): For each m-BRS $(Q^0, \ldots, Q^m)$, there exists a type structure $\mathcal{T}$ so that

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

However, this conjecture is incorrect. The next series of examples will illustrate the issues involved.

Counterexamples It will be useful to begin by showing that Equation (2) cannot be strengthened to require equality.

Example 6.1. Consider the game in Figure 6.1. Let (Q^0, Q^1, Q^2) be the decreasing sequence of product sets with

$$Q_a^1 \times Q_b^1 = S_a \times \{y_1 q_1, y_1 q_2, y_2\} \qquad \text{and} \qquad Q_a^2 \times Q_b^2 = \{x_2\} \times Q_b^1.$$

Note this is a 2-BRS.[10] But, we show that there is no type structure $\mathcal{T}$ with $Q^1 \subseteq \mathrm{proj}_S R^1(\mathcal{T})$ and $Q^2 = \mathrm{proj}_S R^2(\mathcal{T})$.

Suppose otherwise. Then, there exists a type t_a so that $(x_1 z_1, t_a) \in R_a^1(\mathcal{T})$. Observe that, at each information set, $\beta_a(t_a)$ must assign probability 1 to $\{y_2\} \times T_b$. But, y_2 is a sequential best response under

[9]Note the following implication: The set of strategies that survives m rounds of EFR, viz. $\mathrm{EFR}_a^m \times \mathrm{EFR}_b^m$, is consistent with the m-BRS. (Use Battigalli and Siniscalchi, 2002 and Proposition 6.1.)

[10]Let us point to three features of the example: First, $x_1 z_1$ and x_2 are both sequential best responses under a CPS that assigns probability 1 to y_2, $x_1 z_2$ is a unique sequential best response under a CPS that assigns probability 1 to y_3, and x_2 is a unique sequential best response under a CPS that assigns probability 1 to $\{y_1 q_1, y_1 q_2\}$. Second, $y_1 q_1$ (resp. $y_1 q_2$) and y_2 are the *only* strategies that are a sequential best response under a CPS that assigns probability 1 to x_2 at the initial information set and then assigns probability 1 to $x_1 z_2$ (resp. $x_1 z_1$) conditional upon observing x_1. Third, y_2 is a unique sequential best response under a CPS that assigns probability 1 to $x_1 z_2$.

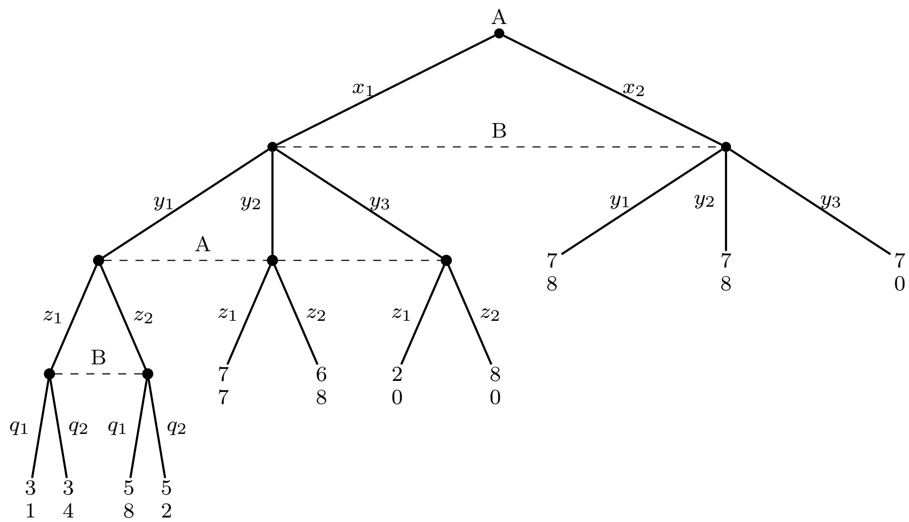

Figure 6.1

every 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, $Q_a^2 \neq \text{proj}_{S_a} R_a^2(\mathcal{T})$. $\square$

Example 6.1 shows that we may have a 2-BRS (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 some $\mathcal{T}$, 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 6.2. *Fix a game* Γ.

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

(ii) For each 2-BRS (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 6.2, Equation (1) does indeed hold for $m = 1, 2$. However, we will next see that an analogue of Proposition 6.2 does not hold for the 3-BRS.

Example 6.2. Return to the game in Figure 6.1. Let (Q^0, Q^1, Q^2, Q^3), where (Q^0, Q^1, Q^2) is the 2-BRS described in Example 6.1 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 on $\{x_1 z_1, x_1 z_2\} \times T_a$. We will argue that $(\{x_1 z_1\} \times T_a) \cap R_a^2(\mathcal{T}) \neq \emptyset$ but $(\{x_1 z_2\} \times T_a) \cap R_a^2(\mathcal{T}) = \emptyset$, 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 6.1 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 $R_b^1(\mathcal{T})$. Thus, $x_1 z_2 \notin \text{proj}_{S_a} R_a^2(\mathcal{T})$. $\square$

Example 6.2 gives a 3-BRS 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 $\text{proj}_S R^3(\mathcal{T})$. The key is that there is a strategy in Q^3_b that is a sequential best response 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 is a strict subset of $\text{proj}_{S_a} R^2_a(\mathcal{T})$.

Let us review what led to the situation where $Q^2_a \subsetneq \text{proj}_{S_a} R^2_a(\mathcal{T})$. The strategy $x_1 z_1$ is a sequential best response under a CPS p_a on S_b. However, for any CPS $\hat{\text{p}}_a$ on $S_b \times T_b$ with $\text{p}_a = \text{marg}_{S_b} \hat{\text{p}}_a$, we have that $\hat{\text{p}}_a$ strongly believes that "Bob is rational." With this in mind, we now restrict attention to a class of games that are generic; in such games, this phenomenon (essentially) cannot arrise.

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 $s_c \in \mathbb{BR}[\text{p}_c]$ if and only if $[s_c] \subseteq \mathbb{BR}[\text{p}_c]$.[11]

Definition 6.4. Call a game **generic** if the following property holds: There exists a CPS $\text{p}_c \in \mathcal{C}(S_{-c}, \mathcal{E}_c)$ so that $s_c \in \mathbb{BR}[\text{p}_c]$ if and only if there exists a CPS $\text{q}_c \in \mathcal{C}(S_{-c}, \mathcal{E}_c)$ so that $[s_c] = \mathbb{BR}[\text{q}_c]$.

Thus, a game is generic if any sequential best response is a "unique" sequential best response under some—perhaps different—CPS. Informally, if a strategy is justifiable then it is "uniquely" justifiable. Here, unique is taken to mean "up to equivalent strategies." A perfect-information game that satisfies no relevant ties in the sense of Battigalli (1997) (Definition D.1) is generic. (See Proposition D.2.) In fact, a game is generic if it satisfies no relevant ties and is degenerately justifiable (Definition D.3), i.e., is a sequential best response under some degenerate CPS. (See Proposition D.1.) A game that satisfies no relevant convexities (Definition 7.1) is also generic. (See Corollary D.1.)

By restricting attention to generic games, we solve the problem that arose in Examples 6.1-6.2. To see this, observe that $x_1 z_1$ is only a sequential best response under a CPS that assigns probability 1 to y_2. But, y_2 is a sequential best response under every CPS. That is, there is no CPS p_b so that y_2 is not a sequential best response under p_b. This occurred despite the fact that $y_1 q_1$ was a sequential best response under some CPS. Genericity requires that, if $y_1 q_1$ is a sequential best response under some CPS, then we can choose the CPS, viz. p^*_b, so that $y_1 q_1$ is the unique sequential best response under p^*_b. (Note, in this game form, $[y_1 q_1] = \{y_1 q_1\}$.) If that were the case, then there would be a CPS under which y_2 is not a sequential best response—namely, p^*_b. As such, we would be able to construct a type structure and a type t_a so that $(x_1 z_1, t_a)$ is rational but t_a *does not* strongly believe that "Bob is rational." (We would require that the type structure have types t_a, t_b with $\text{marg}_{S_a} \beta_b(t_b) = \text{p}^*_b$, and $\beta_a(t_a)((y_2, t_b)|S_b \times T_b) = 1$.) This would solve the problem seen in Examples 6.1-6.2.

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

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

(i) For each type structure $\mathcal{T}$, $(\text{proj}_S R^0(\mathcal{T}), \ldots, \text{proj}_S R^m(\mathcal{T}))$ forms an m-BRS.

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

[11]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.

Part (i) is a special case of Proposition 6.1. Part (ii) is specific to generic games. It says that, for a generic game and an associated m-BRS, 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}. Thus, for generic games, Equation (1) does hold.

Sketch of Proof We provide a sketch of the proof of Theorem 6.1(ii). Fix a generic game Γ and a 2-BRS (Q^0, Q^1, Q^2). The goal is to construct a type structure $\mathcal{T}$ 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.)

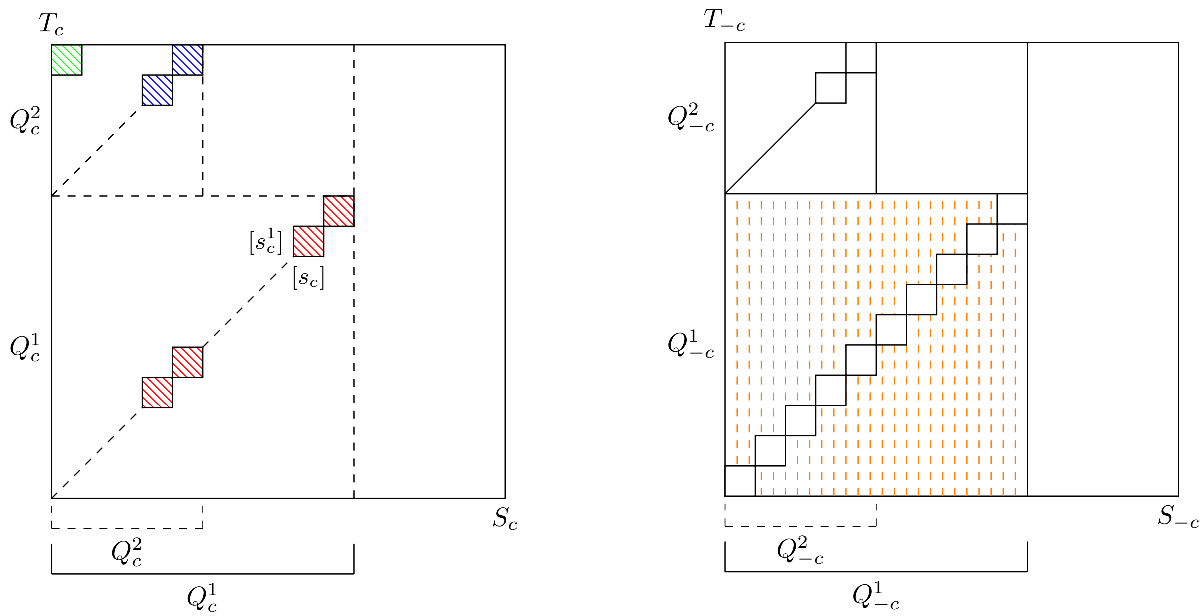

(a) Type Sets

(b) Belief $\beta_c(s_c^1)(\cdot|S_{-c}(h) \times T_{-c})$

Figure 6.2: Construction of Type Structure

Figure 6.2a illustrates the set of strategy-type pairs for player c. The horizontal axis corresponds to the set of strategies; it illustrates $Q_c^2 \subseteq Q_c^1 \subseteq S_c$. The vertical axis corresponds to the set of types that we now construct. In particular, we take $T_c = Q_c^1 \bigsqcup Q_c^2$; that is, T_c is the disjoint union of Q_c^1 and Q_c^2. In doing so, we think of each $s_c \in Q_c^2 \subseteq Q_c^1$ as being associated with two types: a 1-type labeled s_c^1 and a 2-type labeled s_c^2. For each $i = 1, 2$ and $s_c \in Q_c^i$, we refer to (s_c, s_c^i) as an i-strategy-type pair. We will be interested in a modified notion of the diagonal of $Q_c^i \times Q_c^i$—one that accounts for equivalent strategies. So, we will think of the **diagonal of** $Q_c^i \times Q_c^i$ as

$$\text{diag}_c^i = \bigcup_{s_c \in Q_c^i} \Big([s_c] \times [s_c^i] \Big).$$

In Figure 6.2a, the diagonal of $Q_c^1 \times Q_c^1$ is the union over red boxes along the pictorial-diagonal of $Q_c^1 \times Q_c^1$ and the diagonal of $Q_c^2 \times Q_c^2$ is the union over blue boxes along the pictorial-diagonal of $Q_c^2 \times Q_c^2$. The **off-diagonal of** $Q_c^1 \times Q_c^1$ is the white area in $Q_c^1 \times Q_c^1$ (formally, $(Q_c^1 \times Q_c^1) \backslash \text{diag}_c^i$).

The idea will be to construct belief maps so that R_c^1 is contained in the union of squares $(Q_c^1 \times Q_c^1) \cup (Q_c^2 \times Q_c^2)$ and R_c^2 is contained in the square $(Q_c^2 \times Q_c^2)$. Moreover, the belief maps will separate 1-types and 2-types based on whether (or not) they strongly believe rationality. Specifically:

(1) If $(s_c, s_c^1) \in Q_c^1 \times Q_c^1$, then (s_c, s_c^1) is rational and does not strongly believe rationality.

(2) If $(s_c, s_c^2) \in Q_c^2 \times Q_c^2$, then $(s_c, s_c^2) \in Q_c^2 \times Q_c^2$ is rational and strongly believes rationality.

Since $(s_c, s_c^i) \in R_c^1$ implies $[s_c] \times \{s_c^i\} \subseteq R_c^1$, these properties of belief maps will give

$$\mathrm{diag}_c^2 \subseteq R_c^2 \subseteq (Q_c^2 \times Q_c^2) \qquad \text{and} \qquad \mathrm{diag}_c^1 \subseteq R_c^1 \backslash R_c^2 \subseteq (Q_c^1 \times Q_c^1).$$

We may well have $\mathrm{diag}_c^2 \subsetneq R_c^2 \subseteq (Q_c^2 \times Q_c^2)$; pictorially, R_c^2 may well contain both the diagonal blue boxes and the (off-diagonal) green box. However, we will require that $\mathrm{diag}_c^1 = R_c^1 \backslash R_c^2 \subseteq (Q_c^1 \times Q_c^1)$. With this, each off-diagonal point in $Q_c^1 \times Q_c^1$ will be irrational. The role of this requirement will become clear below.

First, we construct the beliefs associated with 2-types. By definition of a 2-BRS, 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 $\mathrm{marg}_{S_{-c}} \beta_c(s_c^2) = j_c(s_c^2)$. Moreover, if $S_{-c}(h) \cap Q_{-c}^1 \neq \emptyset$, we require that $\beta_c(s_c^2)(\cdot | S_{-c}(h) \times T_{-c})$ be concentrated on the diagonal of $Q_{-c}^1 \times Q_{-c}^1$. (We can do this since, in that case, $j_c(s_c^2)(Q_{-c}^1 | S_{-c}(h)) = 1$.)

Next construct the beliefs associated with 1-types. Since the game is generic, for each $s_c \in Q_c^1$, there is a CPS $j_c(s_c^1)$ on $(S_{-c}, \mathcal{E}_c)$ so that $[s_c] = \mathbb{BR}[j_c(s_c^1)]$. For the purpose of illustrating the construction, suppose that Q_{-c}^1 has at least two non-equivalent strategies.[12] Figure 6.2b illustrates this case; note that the off-diagonal (illustrated by the orange lines) is non-empty. Moreover, the off-diagonal meets each strategy in Q_{-c}^1. (Formally, for each $s_{-c} \in Q_{-c}^1$, $(\{s_{-c}\} \times T_{-c}) \cap ((Q_{-c}^1 \times Q_{-c}^1) \backslash \mathrm{diag}_{-c}^1) \neq \emptyset$.) We can then choose $\beta_c(s_c^1)$ so that (i) $\mathrm{marg}_{S_{-c}} \beta_c(s_c^1) = j_c(s_c^1)$, (ii) for each h, $\beta_c(s_c^1)(S_{-c} \times Q_{-c}^1 | S_{-c}(h) \times T_{-c}) = 1$, and (iii) for each h, $\beta_c(s_c^1)(\mathrm{diag}_{-c}^1 | S_{-c}(h) \times T_{-c}) = 0$. So, each $\beta_c(s_c^1)$ has beliefs that are concentrated on 1-states, but off the diagonal.

Observe that, under the construction,

$$R_c^1 = \bigcup_{i=1,2} \bigcup_{s_c^i \in Q_c^i} (\mathbb{BR}[j_c(s_c^i)] \times \{s_c^i\}) = \mathrm{diag}_c^1 \cup \bigcup_{s_c^2 \in Q_c^2} (\mathbb{BR}[j_c(s_c^2)] \times \{s_c^2\}).$$

Since the same holds for the other player, viz. $-c$, the off-diagonal points of $Q_{-c}^1 \times Q_{-c}^1$ are irrational and the diagonal points of $Q_{-c}^1 \times Q_{-c}^1$ are rational. Thus, each 1-type s_c^1 does not strongly believe R_{-c}^1 while each 2-type s_c^1 strongly believes R_{-c}^1. As such,

$$R_c^2 = \bigcup_{s_c^2 \in Q_c^2} (\mathbb{BR}[j_c(s_c^2)] \times \{s_c^2\}).$$

From this it follows that $\mathrm{proj}_{S_c} R_c^1 = Q_c^1$ and $\mathrm{proj}_{S_c} R_c^2 = Q_c^2$: By construction, $Q_c^1 \subseteq \mathrm{proj}_{S_c} R_c^1$ and $Q_c^2 \subseteq \mathrm{proj}_{S_c} R_c^2$. Moreover, each $\mathbb{BR}[j_c(s_c^2)] \subseteq Q_c^2 \subseteq Q_c^1$. So, $\mathrm{proj}_{S_c} R_c^1 \subseteq Q_c^1$ and $\mathrm{proj}_{S_c} R_c^2 \subseteq Q_c^2$.

7 Computational Issues

Refer to Section 4 and Theorem 6.1: In a generic games, we can determine that a strategy has an identified rationality bound of m (with $\infty > m \geq 1$), provided we can compute all the sets consistent with both the m-BRS and the $(m-1)$-BRS.

[12] The proof treats the case of $Q_{-c}^1 = [s_{-c}]$ differently. There, by genericity, we can choose $j_c(s_c^1)$ so that it does not

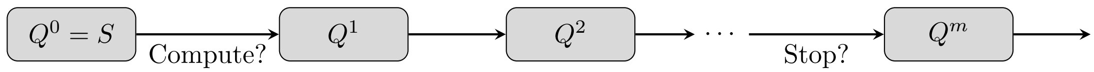

Figure 7.1: m-BRS Elimination Procedure

From a computational perspective, there are two obstacles that hinder implementing the procedure. Refer to Figure 7.1: First, to determine if (Q^0, Q^1) is an 1-BRS, we need to compute the sets $\mathbb{BR}[\mathrm{p}_c]$ for all CPS's p_c. Analogously, for any higher-order m-BRS. However, there are uncountably many such CPS's. Second, we must determine that the procedure stops. Because the game is finite, there must exist some M so that the m-BRS stops shrinking at M, i.e., $Q^M = Q^m$ for all $m \geq M$. However, from the perspective of implementing the procedure, the researcher must know *when* it stops shrinking. We will see that this step is not obvious. This section addresses both computational issues.

Preliminary Step: Simplifying the m-BRS To address the first computational issue, we begin with a useful preliminary result. Toward that end, consider the following thought exercise: The analyst seeks to check whether a particular (Q^0, Q^1) is a 1-BRS. Fix $s_c \in Q^1_c$ and imagine we find some array p_c so that $s_c \in \mathbb{BR}[\mathrm{p}_c]$; that is, imagine we find a sequence of measures—one for each information set—so that s_c is a sequential best response under p_c. The sequence may not be a CPS; that is, it may not satisfy the rules of conditional probability. However, we can always convert the array to a CPS, viz. q_c, so that s_c remains a sequential best response under q_c. In fact, this can be done in a way that preserves strong belief. The issue is that there may be new strategies that become a sequential best response. Thus, even if the array satisfies the maximality property, the CPS may not. (See Example D.1, which is a non-generic game.)

This raises the question: Are there situations in which we can ensure that the CPS satisfies the maximality criterion, when the array does so? If so, we can simplify the definition of an m-BRS by only making reference to an array. If not, we must ensure that we construct an appropriate CPS.

We will now see that, in a particular class of generic games, we can indeed ensure that the CPS satisfies the maximality criterion, whenever the array does so. The key is that, in that class of games, we can construct the CPS q_c so that $\mathbb{BR}[\mathrm{q}_c] = [s_c]$. Note, since $s_c \in \mathbb{BR}[\mathrm{p}_c]$, $[s_c] \subseteq \mathbb{BR}[\mathrm{p}_c]$. (That is, all strategies that are equivalent to s_c must be a sequential best response under the array p_c.) So, if p_c satisfies the maximality criterion, so does q_c. This ensures that, in that class of games, it suffices to restrict attention to arrays.

Fix some $X_{-c} \subseteq S_{-c}$ and some information set $h \in H_c$ with $s_c \in S_c(h)$. Say r_c **supports** s_c **with respect to** (X_{-c}, h) if there exists $\sigma \in \mathcal{P}(S_c(h))$ with (1) $\sigma(r_c) > 0$, and (2) for all $s_{-c} \in X_{-c} \cap S_{-c}(h)$, $\sum_{r_c \in S_c(h)} \pi_c(r_c, s_{-c})\sigma(r_c) = \pi_c(s_c, s_{-c})$. (Note, if $s_c \notin S_c(h)$, then no r_c supports s_c with respect to (X_{-c}, h).)

Definition 7.1. The game satisfies **no relevant convexities** (**NRC**) if, for each $h \in H_c$, the following holds: If $s_c \in S_c(h)$ and r_c supports s_c with respect to some (X_{-c}, h), then $\zeta(s_c, s_{-c}) = \zeta(r_c, s_{-c})$ for each $s_{-c} \in X_{-c} \cap S_{-c}(h)$.

Informally, a game satisfies NRC if strategies in the support of a mixture $\sigma \in \mathcal{P}(S_c)$ induce the same path of play as s_c, whenever player c is indifferent between σ and s_c. As stated earlier, games that satisfy NRC are generic. (This is Corollary D.1.)

strongly believe Q^1_{-c}.

Proposition 7.1. *Fix a game that satisfies NRC. Then* $(Q^0, \ldots, Q^m)$ *forms an* m*-BRS if and only if* Q^1 *is non-empty and, for each* $n = 1, \ldots, m$*, the following hold:*

(i) For each $s_c \in Q_c^n$*, there exists an array* $\mathrm{p}_c \in \mathcal{A}(S_{-c}, \mathcal{E}_c)$ *that strongly believes* $Q_{-c}^0, \ldots, Q_{-c}^{n-1}$*, so that* $s_c \in \mathbb{BR}[\mathrm{p}_c]$*.*

(ii) A strategy $s_c \in Q_c^n$ *if and only if* $[s_c] \subseteq Q_c^n$*.*

Fix a game that satisfies NRC and a decreasing sequence of product sets $(Q^0, \ldots, Q^{m-1}, Q^m)$ so that $(Q^0, \ldots, Q^{m-1})$ is an $(m-1)$-BRS. We seek to determine whether $(Q^0, \ldots, Q^m)$ is also an m-BRS. Proposition 7.1 provides two ways that simplify making this determination. First, we can replace CPS's with arrays. Second, we can replace the maximality criterion with a requirement that if $s_c \in Q_c^m$ then Q_c^m includes all equivalent strategies. (Unlike the maximality criterion, this does not make reference to arrays.)

Problem 1: Computing Best Responses Fix an m-BRS, viz. $(Q^0, \ldots, Q^m)$, and some $Q = Q_a \times Q_b \subseteq Q^m$. We seek to determine whether or not this forms an $(m+1)$-BRS. Toward that end, we would need to search across all CPS's p_c, asking that they satisfy conditions (BRP.1)-(BRP.2)-(BRP.3) of an $(m+1)$-BRS. The set of such CPS's is uncountable.

But, we have seen that, when the game satisfies NRC, there is an alternate approach based on Proposition 7.1: We will think of a strategy $s_c \in Q_c^m$ as "passing the test" if there is an array $\mathrm{p}_c \in \mathcal{A}(S_{-c}, \mathcal{E}_c)$ so that $s_c \in \mathbb{BR}[\mathrm{p}_c]$ and p_c strongly believes $Q_{-c}^0, \ldots, Q_{-c}^m$. So, when the game satisfies NRC, it suffices to check that Q_c can be written as the union over sets $[s_c] \subseteq Q_c^m$ so that s_c passes the test.

This alternative approach has two pros and one con. First, it forgoes checking condition (BRP.3) of an m-BRS. Second, if the researcher locates an array that does satisfy conditions (BRP.1)-(BRP.2) of an m-BRS, the researcher does not have to verify that the array is a CPS. But, third, to show that a strategy fails the test, we now have to quantify across all arrays. The key is that this can be achieved by making use of the simplex algorithm. We now explain how it can be used.

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

$$Q_{-c}^{n(h)} \cap S_{-c}(h) = \{s_{-c}^1, \ldots, s_{-c}^K\} \qquad \text{and} \qquad S_c(h) = \{s_c^1, \ldots, s_c^L\}.$$

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 non-negative numbers $\mu^1, \ldots, \mu^K$ with $\sum_{k=1}^K \mu^k = 1$, so that s_c maximizes $\sum_{k=1}^K \pi_c(\cdot, s_{-c}^k)\mu^k$ among all strategies in $S_c(h) = \{s_c^1, \ldots, s_c^L\}$. A strategy s_c **passes the test** if it passes the test at each $h \in H_c$.

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 $(\mu^1, \ldots, \mu^K, \tau^1, \ldots, \tau^L)$ to solve

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

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.

Remark 7.1. This alternate approach can resolve a background computational issue: When the strategy set is large, the set of all possible subsets of Q_c of S_c may be large. Thus, it may be difficult to test all sequences (Q^0, Q^1) to determine if they form a 1-BRS.

When the game satisfies NRC, we can simplify this step. Instead of testing each subset of strategies, we can test the strategies themselves. The sequence (Q^0, Q^1) is a 1-BRS if and only if, for each c, Q^1_c can be written as a union of sets $[s_c]$ where s_c passes the test just described. If a computer program is used to implement the simplex algorithm, the program can be written to compute these unions. (We thank a referee for raising this issue.) □

Problem 2: Termination of the Procedure Fix a decreasing sequence of strategies $(Q^0, Q^1, Q^2, \ldots)$, where each $(Q^0, \ldots, Q^m)$ forms an m-BRS. Since each $Q^{m+1} \subseteq Q^m$, $(Q^0, Q^1, Q^2, \ldots)$ defines an iterative elimination procedure. We refer to this as an m-BRS elimination procedure. Note, there may be many such elimination procedures, corresponding to distinct $(Q^0, Q^1, Q^2, \ldots)$ and $(\hat{Q}^0, \hat{Q}^1, \hat{Q}^2, \ldots)$.

Because the strategy set is finite, this elimination 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, he could use that fact to determine that the elimination procedure has stopped.

At first glance, there may appear to be straightforward route to determine 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-BRS elimination procedure. We may have $Q^{m+1} \subsetneq Q^m = Q^{m-1}$.

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

Figure 7.2

Example 7.1. Consider the simultaneous-move game given by Figure 7.2 and note that the game is generic. Yet, for each m, there is an m-BRS with $(Q^0, \ldots, Q^m)$, so that (i) for each $n \leq m$, $Q^n = \{U, D\} \times \{L, R\}$, and (ii) $Q^{m+1} = \{U, D\} \times \{R\}$. Thus, the $(m+1)$-BRS procedure has no shrinkage up until round m, but a shrinkage at round $(m+1)$. □

To understand why Example 7.1 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 7.1 highlights the fact that we can have arbitrarily long pauses before shrinkage: For any M, there is some m-BRS procedure $(Q^0, Q^1, Q^2, \ldots)$ so that the procedure has not terminated within M steps. (That is, we can have $Q^0 = Q^1 = \cdots = Q^M$ and $Q^{M+1} \subsetneq Q^M$.)

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-BRS procedure $(Q^0, Q^1, Q^2, \ldots)$ with a pause at round m, i.e., $Q^{m+1} = Q^m$ but $Q^{m+2} \subsetneq Q^{m+1}$. The key is that any eliminated strategy—i.e., any strategy in $Q^{m+1}_c \backslash Q^{m+2}_c$—must be contained in $\overline{S}^{m+2}$. That is, there must exist some other m-BRS procedure $(\hat{Q}^0, \hat{Q}^1, \hat{Q}^2, \ldots)$ so that $Q^{m+1} \backslash Q^{m+2} \subseteq \hat{Q}^{m+2}$. This follows from the following:

Observation 7.1. *Fix some* $(Q^0, Q^1, Q^2, \ldots)$ *where, for each* m, $(Q^0, \ldots, Q^m)$ *is an* m*-BRS. 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 $\hat{Q}^n = Q^m$, for each $n \geq m$. As such, $Q^m \subseteq \overline{S}^n$ for all $n \geq m$. From this, the following termination result follows:

Proposition 7.2. *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 7.2 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 $\overline{M}$-BRSs, $(Q^0, \ldots, Q^{\overline{M}})$.

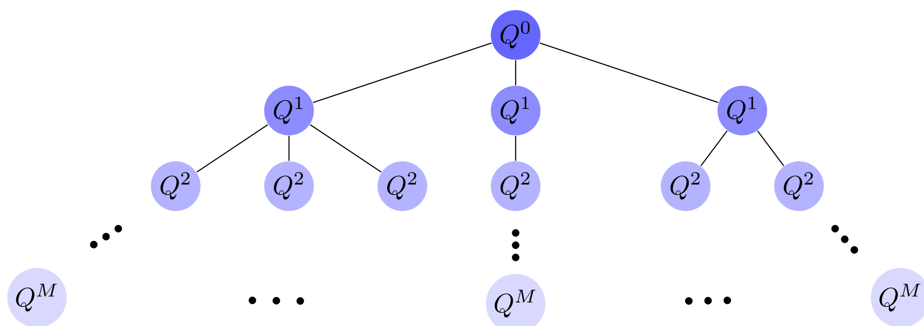

Figure 7.3: m-BRS Elimination Tree

In practice, it is often not necessary to compute all the $\overline{M}$-BRSs. Refer to Figure 7.3. Begin with $Q^0 = S$ and identify all the 1-BRSs (Q^0, Q^1). Use these 1-BRS's to identify all the 2-BRSs (Q^0, Q^1, Q^2). Etc. We can stop after we have identified all the $\overline{M}$-BRSs. However, along any given $\overline{M}$-BRS path $(Q^0, Q^1, \ldots, Q^{\overline{M}})$, it may be possible to stop prior to round $\overline{M}$. In fact, we can stop at $m < \overline{M}$ if $Q^m = Q^{m+1}$.

8 Non-Trivial Identification

Return to the BoS with an Outside Option. Laboratory experiments have shown that a significant fraction of subjects do play *Out*. (See Cooper, DeJong, Forsythe and Ross 1993 and Brandts and Holt 1995.) This is interpreted as evidence that players are not fully strategically sophisticated (in the sense of engaging in forward-induction reasoning). But, we have seen that this behavior is, in fact, consistent with RCSBR.

That said, BoS with an Outside Option is not suitable for the purpose of identifying the rationality bound. This is because the partitions $\mathcal{IB}_a$ and $\mathcal{IB}_b$ are too coarse. Specifically,

$$\mathcal{IB}_a = \{\text{IB}_a^0, \text{IB}_a^\infty\} \qquad \text{and} \qquad \mathcal{IB}_b = \{\text{IB}_b^\infty\},$$

where $\text{IB}_a^0 = \{\textit{In-D}\}$, $\text{IB}_a^\infty = \{\textit{Out}, \textit{In-U}\}$, and $\text{IB}_b^\infty = \{L, R\}$. Thus, the most the researcher can hope to learn is (i) whether the subject reveals herself to be irrational, or (ii) whether that the subject's behavior is consistent with RCSBR.

With this in mind, say that the game **displays trivial identification for c** if either $\mathcal{IB}_c = \{\mathrm{IB}_c^0, \mathrm{IB}_c^\infty\}$ or $\mathcal{IB}_c = \{\mathrm{IB}_c^\infty\}$. There is a class of games studied in the experimental literature which display trivial identification for both players. Notable examples include the burning money game (Huck and Müller, 2005), Bagwell-Ramey (Bagwell and Ramey, 1996) style entry deterrence (Brandts, Cabrales and Charness, 2007), and BoS with an outside coordination game (Brandts and Holt, 1995). These games are arguably not suitable for the purpose of identifying the rationality bound. This raises the question: Are there games which permit *non-trivial* identification for players? Here, we show that the Centipede game provides non-trivial identification for the first mover.[13]

Centipede Game Figure 8.1 depicts the Centipede game. We order the non-terminal nodes (or vertices) as $v = 1, 2, \ldots, V$, where $V \geq 3$. (So, $v = 1$ indicates the initial node and $v = V$ indicates the last non-terminal node.) If the game ends after out_v is played and v is odd (resp. even), then the payoffs are $(x + (v-1)y, x + (v-2)y)$ (resp. $(x + (v-3)y, x + vy)$), where $x, y > 0$.[14] If the game ends after in_V is played and V is odd (resp. even), then the payoffs are $(x+(V-2)y, x+(V+1)y)$ (resp. $(x+Vy, x+(V-1)y)$). Figure 8.1 depicts V odd.

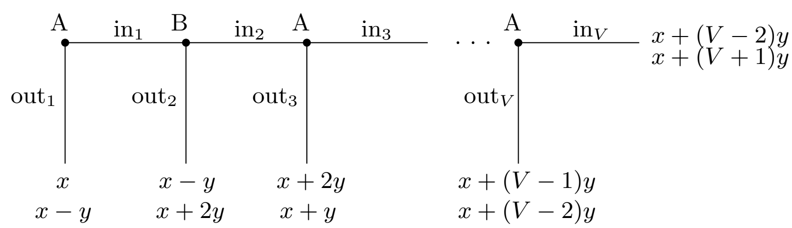

Figure 8.1: Centipede Game

Write ℓ for the last player and $-\ell$ for the second-to-last player. If V is odd then ℓ is Ann and $-\ell$ is Bob; if V is even then ℓ is Bob and $-\ell$ is Ann. Let $[\mathrm{out}, v]_c$ be the set of strategies of player c that allow v and then play out_v. Likewise, let $[\mathrm{in}]_c$ be the set that contains the (unique) strategy of player c that specifies in_v at every node v.

Let us point to four (related) features of the game: First, the game is generic (Proposition D.2). Thus, we can exploit Theorem 6.1. Second, 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}. Third, the player who moves at vertex V, strictly prefers out_V to in_V. Fourth, for each $v \leq V - 1$ (resp. $v = V$), $[\mathrm{out}, v]_c$ is the set of best responses under a CPS that strongly believes $[\mathrm{out}, v+1]_{-c}$ (resp. $[\mathrm{in}]_{-c}$).[15]

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. Observe that

$$\mathrm{EFR}_\ell^1 \times \mathrm{EFR}_{-\ell}^1 = (S_\ell \backslash [\mathrm{in}]_\ell) \times S_{-\ell} \qquad \text{and} \qquad \mathrm{EFR}_\ell^2 \times \mathrm{EFR}_{-\ell}^2 = \mathrm{EFR}_\ell^1 \times (S_{-\ell} \backslash [\mathrm{in}]_{-\ell}).$$

[13] We use the phrase non-trivial identification for c to reflect a situation where $\mathrm{IB}_c^0 \cup \mathrm{IB}_c^\infty \neq S_c$. Thus, there may be degrees of non-trivial identification. In one game, we may have non-trivial identification with $\mathrm{IB}_c^0 \cup \mathrm{IB}_c^1 \cup \mathrm{IB}_c^\infty = S_c$, whereas in a second game we may have $\mathrm{IB}_c^0 \cup \mathrm{IB}_c^1 \cup \mathrm{IB}_c^\infty \neq S_c$. In principle, the latter game can provide the researcher with more information about the players' RmSBR bound.

[14] The first component in the payoff vector is Ann's payoffs.

[15] Reny's TOL game also satisfies these properties. Thus, anything we say about Centipede also applies to that game.

Moreover,

$$\mathrm{EFR}_\ell^m \times \mathrm{EFR}_{-\ell}^m = \begin{cases} (\mathrm{EFR}_\ell^{m-1} \backslash [\text{out}, V+3-m]_\ell) \times \mathrm{EFR}_{-\ell}^{m-1} & \text{if } m = 3, \ldots, V \text{ is odd,} \\ \mathrm{EFR}_\ell^{m-1} \times (\mathrm{EFR}_{-\ell}^{m-1} \backslash [\text{out}, V+3-m]_{-\ell}) & \text{if } m = 4, \ldots, V \text{ is even.} \end{cases}$$

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

Unlike EFR, the m-BRS procedure has very different implications for the first mover (Ann) and the second mover (Bob).

Proposition 8.1. *In the Centipede game, the following hold for each finite $m \geq 1$:*

(i) $\overline{S}_a^m = \mathrm{EFR}_a^m$.

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

Proposition 8.1 implies that

$$\mathcal{IB}_b = \begin{cases} \{\mathrm{IB}_b^\infty\} & \text{if } V \text{ is odd,} \\ \{\mathrm{IB}_b^0, \mathrm{IB}_b^\infty\} & \text{if } V \text{ is even.} \end{cases}$$

Thus, the game displays trivial identification for Bob. However, this is not the case for Ann. For instance, when V is odd,

$$\mathcal{IB}_a = \{\mathrm{IB}_a^0, \mathrm{IB}_a^2, \mathrm{IB}_a^4, \ldots, \mathrm{IB}_a^{V-1}, \mathrm{IB}_a^\infty\},$$

where $\mathrm{IB}_a^0 = [\text{in}]_a$, $\mathrm{IB}_a^\infty = [\text{out}, 1]_a$, and, for each $2m = 2, \ldots, V-1$, $\mathrm{IB}_a^{2m} = [\text{out}, V-2m+2]_a$.

Why does a difference arise between the first mover (Ann) and the second mover (Bob)? To address this question, focus on the case where V is odd. There, $[\text{out}, 1]_a \times S_b$ is an EFBRS. Thus, $S_b \subseteq \overline{S}_b^\infty$. But, for any non-empty EFBRS, viz. $Q_a \times Q_b$, we have $Q_a = [\text{out}, 1]_a$.[16] Thus, $\overline{S}_a^\infty = [\text{out}, 1]_a$.

At first glance, part (i) of Proposition 8.1 may appear trivial: For each m, $\mathrm{EFR}_a^m \times \mathrm{EFR}_b^m$ is consistent with the m-BRS. Thus, $\mathrm{EFR}_a^m \subseteq \overline{S}_a^m$. However, the key is showing that $\overline{S}_a^m \subseteq \mathrm{EFR}_a^m$ and, as we have seen, this is not the case for the second mover Bob. (Appendix F explains why this is the case.)

The implications for identifying the rationality bound are quite stark. Centipede displays trivial identification for the second mover. But, that is not the case for the first mover. Thus, first-mover (and only first-mover) Centipede data is suitable for identifying the rationality bound. Section 9.g provides a *preliminary* discussion of possibilities for tighter identification.

9 Discussion

9.a Limited Ability to Engage in Interactive Reasoning This paper focuses on bounded reasoning about rationality. As discussed in the Introduction (page 5), we take no position on *why* a player may exhibit such rationality bounds. Ann may face rationality 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

[16]Suppose not. Then there is some $(s_a, s_b) \in Q_a \times Q_b$ where Ann plays in_1. Consider the strategy profile that results in the maximum path of play of in: specifically, it results in $\text{in}_1 \cdots \text{in}_v$ and the player, viz. $-c$, who moves at $v+1$ plays out_{v+1}. Let s_c be a strategy in Q_c that plays in up to and including v. Any CPS that strongly believes Q_{-c} must, at v, assign probability 1 to $-c$ playing out_{v+1}. Thus, s_c cannot be a sequential best response under any CPS that strongly believes Q_{-c}.

behavior. Or, alternatively, the rationality bounds may be an artifact of a limited ability to engage in interactive reasoning—i.e., limits on the ability to compose sentences of the form "I think that you think that I think ...."

At first glance, the latter scenario may seem inconsistent with our epistemic framework: Type structures induce infinite hierarchies of beliefs about the strategies played—i.e., m^{th}-order beliefs for all m. This suggests that players can compose sentences of the form "I think that you think that I think ...." However, this inconsistency is illusory. The key observation is that, if Ann has an identified rationality bound of m, then hierarchies of beliefs beyond level m do not affect the behavioral characterization of R$(m-1)$SBR. Formally, consider two types t_a and u_a with the same m^{th}-order beliefs about the strategies played. For any strategy s_a, the strategy-type pair (s_a, t_a) is consistent with R$(m-1)$SBR if and only if (s_a, u_a) is consistent with R$(m-1)$SBR. The higher-order beliefs become an artifact of our formalism and do not have implications for our characterization result. So, in particular, we could instead adapt the type structure frameworks in Kets (2010) and Heifetz and Kets (2018)—which allow for finite-order beliefs—to conditional beliefs and appropriately apply R$(m-1)$SBR in that framework. We would reach analogous conclusions. (Appendix A in Heifetz and Kets, 2018 makes a similar point, in a different context.)

9.b Portability of the Rationality Bound While we take no position on *why* a player may exhibit rationality bounds, the source can have implications for the portability of the *identified* rationality bound—i.e., what the identified rationality bound would be when we change the game or the subject pool. To illustrate why, first, consider the case where Ann's rationality bound is not determined by limits on the ability to engage in interactive reasoning.[17] Suppose, in a given game, Ann has an actual rationality bound of m—i.e., she is characterized by some $(s_a, t_a) \in R_a^m(\mathcal{T})$, where t_a does not strongly believe $R_b^m(\mathcal{T})$. Since she is capable of more than m steps of "I think, you think, I think,...," *ex ante*, she may well assign positive probability to $R_b^m(\mathcal{T})$. (She is capable of doing so.) More informally, she may have non-degenerate beliefs on reasoning about rationality. Such non-degenerate beliefs can cause the identified rationality bound to vary across games even if the actual rationality bound does not. (Friedenberg, Kets and Kneeland 2016 make this point and, in the context of a particular experiment, show that such non-degenerate beliefs about rationality explain the subjects' behavior.) Second, consider the case where Ann's rationality bound is determined by limits on the ability to engage in interactive reasoning. In that case, her identified rationality bound will coincide with her actual rationality bound and the bound on her ability to engage in interactive reasoning. Thus, her identified rationality bound can only vary (i.e., across games or subject pools), if her ability bound also varies. The fact that this ability bound may vary is suggested by Alaoui and Penta's (2016) model of endogenous depth of reasoning.

9.c Strong Belief versus Initial Belief Example 3.1 explained why we focus on "strong belief of rationality" and not "full belief of rationality." An alternative 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 1 to the event that the other player is rational.

Initial belief relaxes what it means to reason about rationality throughout the game. (It does not require a player to rationalize past behavior, even if it is possible to do so.) As such, it allows the researcher to rationalize the data at higher levels: If at a given state there is RmSBR, then there is also rationality and

[17] The possibility that this may occur was first suggested (at least, in print) by Agranov, Potamites, Schotter and Tergiman (2012) in the context of level-k models. See Alaoui and Penta (2017) (level-k) and Friedenberg, Kets and Kneeland 2016 (RmBR) for more on this.

m^{th}-order initial belief of rationality. Thus, if the data has an identified rationality bound of m according to "strong belief," then it has an identified rationality bound of $n \geq m$ according to "initial belief."

The implication is that, under initial belief, it is more difficult to identify levels of reasoning about rationality. For instance, refer back to the three-legged Centipede game. There, both $[\text{out}]_1$ and $[\text{out}]_3$ are consistent with rationality and common initial belief of rationality. In fact, this is true more generally: In Centipede, initial belief would give rise to trivial identification for all players. Thus, the first player's behavior would not be suitable for the purpose of identifying a rationality bound based on initial belief.

9.d 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 opponents? (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. (Take $-c$ to mean all players $i \neq c$.)

The question of correlation vs. independence has important implications for trivial vs. non-trivial identification. As an example, take Dufwenberg and Van Essen's (2018) King of the Hill (line) experiment: Under correlated rationalization, there is trivial identification for the second mover; but, under independent rationalization, there is non-trivial identification for the second mover.

9.e Identified Bound of ∞ Refer back to Remark 4.2. We defined a strategy to have an identified rationality bound of ∞ if it is consistent with RmSBR for all m. Alternatively, one might have defined a strategy to have an identified rationality bound of ∞ if it is consistent with RCSBR. In light of Observation 7.1, we can conclude that the two definitions are 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 $s_c \in \text{proj}_{S_c} \bigcap_m R_c^m(\mathcal{T})$.

9.f Coarser Datasets We have assumed that the researcher observes a dataset that is a subset of the strategies. This is best thought of as having the data arise from an experiment that employs the strategy method. But, often, the researcher has access to a dataset where observations are a signal of the strategies played. For instance, the researcher may have access to experimental data that uses the direct response method; in that case, the dataset consists of observed paths of play or, equivalently, terminal nodes. The same would obtain if the researcher had access to a path of bids in an auction. If the researcher instead had access to auction data that describes the winner and the price paid, the dataset consists of outcomes of the game. (Each path of play induces a unique terminal node; many terminal nodes can be associated with one outcome.)

In each of these cases, there is a dataset $\mathbb{D}$ and a mapping $\delta : S_a \times S_b \to \mathbb{D}$. Say the data $d \in \mathbb{D}$ is **consistent with RmSBR for** c if there exists a strategy s_c consistent with RmSBR so that $(\{s_c\} \times S_{-c}) \cap \delta^{-1}(\{d\}) \neq \emptyset$. In words, the data is consistent with RmSBR if there exists a strategy s_c that is consistent with both the data and R$(m-1)$SBR. Then, analogous to Definition 4.1:

Definition 9.1.

(1) Say $d \in \mathbb{D}$ has a **c-identified rationality bound of** 0 if the data is inconsistent with R0SBR for c.

(2) For finite $m \geq 1$, say $d \in \mathbb{D}$ has a **c-identified rationality bound of** m if the data is consistent with R$(m-1)$SBR for c, but is inconsistent with RmSBR for c.

(3) Say $d \in \mathbb{D}$ has a **c-identified rationality bound of** ∞ if, for each finite $m \geq 1$, the data is consistent with RmSBR for c.

On the surface, this setup appears different from Section 4. Here, the strategy method corresponds to a dataset $\mathbb{D} \subseteq S_a \times S_b$ and the identity map δ. With this, the identification in Definition 9.1 corresponds to Definition 4.1. When the dataset $\mathbb{D}$ is coarser—e.g., if it is a subset of terminal nodes or outcomes—then there are typically many strategies consistent with the data. As such, it becomes easier to rationalize the data at higher levels and, so, non-trivial identification becomes more difficult to achieve.

As an illustration, return to Centipede and take V to be odd. Consider a subject who, in the role of Ann, would choose a strategy in $[\text{out}, V]_a$. If the researcher employs the strategy method, he would observe the strategy and the identified rationality bound would be 1. If, however, the researcher employs the direct response method and the subject is matched with someone who plays out_2 in the role of Bob, then the identified rationality bound would be V; this is because the researcher only observes that the subject chose in_1 and that choice is consistent with strategies in $[\text{out}, 3]_a$. While this problem can be significant in theory, it may or may not be significant in practice. For instance, it does not *appear* to be a significant hinderance in the McKelvey and Palfrey (1992) dataset. We discuss this further in Appendix G.

9.g Finer Datasets We have implicitly assumed that the researcher only observes behavior in one game Γ. The researcher, however, may have repeat observations or observations across different games. If so, he may be able to leverage this information to provide tighter identification.

Centipede illustrates how the researcher might leverage such information. When V is odd, the researcher can use first-mover data to separate out an identified rationality bound of $2m$ from an identified rationality bound of $2m+2$, provided $2m+2 < V$. However, in that game, there is no strategy s_a with an identified rationality bound of $2m+1$. Thus, if the first mover's actual rationality bound is $2m+1$, the researcher will identify her rationality bound as at least $2m+2$.[18] That is, there is no data that *could* distinguish an actual rationality bound of $2m+1$ from an actual rationality bound of $2m+2$.

However, suppose the researcher observes the strategies played both in Centipede with V odd and in Centipede with $V' = V+1$ even. In the latter Centipede,

$$\mathcal{IB}'_a = \{\text{IB}^1_a, \text{IB}^3_a, \ldots, \text{IB}^V_a, \text{IB}^\infty_a\},$$

where $\text{IB}^1_a = [\text{in}]_a$, $\text{IB}^\infty_a = [\text{out}, 1]_a$, and, for each $2m = 2, \ldots, V+1$, $\text{IB}^{2m-1}_a = [\text{out}, V - 2m + 4]_a$. If we take the identified rationality bound to be the minimum of the two bounds identified in V and V', then, for each $m < V$, there is some data that would have an identified rationality bound of m. Thus, there is data that *could* distinguish an actual rationality bound of $2m+1$ from an actual rationality bound of $2m+2$, provided $2m+2 < V$.

[18] Refer back to Remark 4.1 on why we say "at least $2m+2$" instead of simply "$2m+2$."

Note, however, in leveraging information across games, the researcher is making an important assumption: If a subject's *actual* rationality bound is m in Centipede with V non-terminal nodes, then it is also m in Centipede with $V' = V + 1$ non-terminal nodes. More formally: Suppose, in the game with V non-terminal nodes, the subject's *actual* behavior and beliefs are characterized by $(s_a, t_a) \in R_a^m(\mathcal{T}) \backslash R_a^{m+1}(\mathcal{T})$. Then, in the game with V' non-terminal nodes, there exists a type structure $\mathcal{T}'$ so that the subject's *actual* behavior and beliefs are characterized by $(s'_a, t'_a) \in R_a^m(\mathcal{T}') \backslash R_a^{m+1}(\mathcal{T}')$.

There is reason to be cautious about this assumption. The literature has pointed out that level-k behavior may vary across games. (See Georganas, Healy and Weber, 2015; Alaoui and Penta, 2016; Cooper, Fatas, Morales and Qi, 2016.) However, in those games, level-k reasoning does not correspond round-for-round with rationalizability; as such, it is not obvious that the identified rationality bound varies across games. Even if the *identified* rationality bound did vary across games, the *actual* rationality bound may not. (See Section 9.b.) To the extent that it may be helpful to make use of behavior across games, the researcher may first want to investigate whether the actual rationality bound varies across games.

9.h Simultaneous-Move Games Proposition 6.2 stated that, if (Q^0, Q^1, Q^2) is a 2-BRS, there exists a type structure $\mathcal{T}$ with $Q^1 = \text{proj}_S R^1(\mathcal{T})$ and $Q^2 \subseteq \text{proj}_S R^2(\mathcal{T})$. Example 6.2 showed that the result does not generalize beyond the 2-BRS. However, in simultaneous-move games, it does: If $(Q^0, \ldots, Q^m)$ is an m-BRS, there is a $\mathcal{T}$ 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-BRS and this set is the set of m-rationalizable strategies.

Why does this difference arise? In simultaneous-move games, strong belief is monotonic; in an extensive form, it is not.[19] Thus, in simultaneous-move games, if a CPS strongly believes Q^2_{-c} and $Q^2_{-c} \subsetneq \text{proj}_{S_{-c}} R^2_{-c}(\mathcal{T})$, it also strongly believes $\text{proj}_{S_{-c}} R^2_{-c}(\mathcal{T})$. However, in an extensive form, 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 it does not strongly believe $\text{proj}_{S_{-c}} R^2_{-c}(\mathcal{T})$.

9.i Beyond Generic Games It would be desirable to have a procedure that determines the sets $\overline{S}^m$ in non-generic games. One possibility would be to amend the definition of an m-BRS. In light of Example 6.2 one might suggest the following: If $s_a \in Q^1_a \backslash Q^2_a$, then there exists some CPS p_a that satisfies conditions (BRP.1)-(BRP.2)-(BRP.3) and does not strongly believe Q^1_b. However, under that amendment, we loose an analogue of Proposition 6.1: For a given $\mathcal{T}$, $(\text{proj}_S R^0(\mathcal{T}), \ldots, \text{proj}_S R^m(\mathcal{T}))$ may fail the new property.

Example 6.2 illustrates that, for a given m-BRS $(Q^0, Q^1, \ldots, Q^m)$, there may be no $\mathcal{T}$ so that $Q^n \subseteq \text{proj}_S R^n(\mathcal{T})$ for each $n = 1, \ldots, m$. The example leaves open that there may be an alternate m-BRS $(\hat{Q}^0, \hat{Q}^1, \ldots, \hat{Q}^m)$ so that the following holds: (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) does hold for all games. We neither know this to be true nor have a counterexample. Thus, we leave it as an open question.

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* $\text{proj}_{S_{-c}} E_{-c}$.

[19] In simultaneous-move games, strong belief coincides with "belief," i.e., *ex ante* assigning probability 1 to an event.

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 $\text{p}_c \in \mathcal{C}(S_{-c}, \mathcal{E}_c)$ and a measurable mapping $\tau_{-c} : S_{-c} \to S_{-c} \times T_{-c}$. Define q_c as follows: For each conditioning event $S_{-c}(h) \times T_{-c} \in \mathcal{E}_c \otimes T_{-c}$ and each Borel $E_{-c} \subseteq S_{-c} \times T_{-c}$, set

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

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 $\text{marg}_{S_{-c}} q_c = p_c$. As a consequence, for any given CPS $\text{p}_c \in \mathcal{C}(S_{-c}, \mathcal{E}_c)$, we can find some CPS $q_c \in \mathcal{C}(S_{-c} \times T_{-c}, \mathcal{E}_c \otimes T_{-c})$ so that $\text{marg}_{S_{-c}} q_c = p_c$.

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 $s_c \in S_c(h) \cap S_c(h')$ 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 $\mu_c \in \mathcal{P}(S_{-c})$ with $\mu_c(S_{-c}(h^*)) = 1$ and $\mu_c(S_{-c}(h^{**})) > 0$. If $s_c \in S_c(h^{**})$ is optimal under μ_c given all strategies in $S_c(h^*)$, then s_c is optimal under $\mu_c(\cdot|S_{-c}(h^{**}))$ given all strategies in $S_c(h^{**})$.*

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

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

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

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

contradicting the hypothesis that s_c is optimal under μ_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 Proofs of Propositions 6.1-6.2

Proof of Proposition 6.1. We will show that, for each $m \geq 1$, $(\text{proj}_S R^0(\mathcal{T}), \ldots, \text{proj}_S R^m(\mathcal{T}))$ forms an m-BRS. 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 $s_c \in \mathbb{BR}[\mathrm{p}_c]$. Moreover, if $r_c \in \mathbb{BR}[\mathrm{p}_c]$, 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, $(\text{proj}_S R^0(\mathcal{T}), \ldots, \text{proj}_S R^m(\mathcal{T}))$ forms an m-BRS. 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 $(s_c, t_c) \in R^1_c(\mathcal{T})$, $s_c \in \mathbb{BR}[\mathrm{p}_c]$. 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 $r_c \in \mathbb{BR}[\mathrm{p}_c]$, then $(r_c, t_c) \in R^{m+1}_c(\mathcal{T})$ and so $r_c \in \text{proj}_{S_c} R^{m+1}_c(\mathcal{T})$. ■

Proof of Proposition 6.2(i). Fix a 1-BRS (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 $[s_c] \in \mathbb{BR}[\mathrm{p}_c] \subseteq Q^1_c$. (The fact that such a CPS exists follows from the definition of a 1-BRS.) It follows that $\text{proj}_{S_c} R^1_c(\mathcal{T}) = Q^1_c$. ■

Proof of Proposition 6.2(ii). Fix a 2-BRS (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 $T_c = Q_c^1$ and define $\beta_c(s_c)$ so that $\text{marg}_{S_{-c}}\beta_c(s_c) = j_c(s_c)$. Moreover, for each h with $S_{-c}(h) \cap Q_{-c}^1 \neq \emptyset$ and each $s_{-c} \in S_{-c}(h) \cap Q_{-c}^1$, set

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

Then,

$$R_c^1(\mathcal{T}) = \bigcup_{s_c \in Q_c^1} (\mathbb{BR}[j_c(s_c)] \times \{s_c\}) \qquad \Longrightarrow \qquad \text{proj}_{S_c} R_c^1(\mathcal{T}) = Q_c^1.$$

Moreover, if $s_c \in Q_c^2$, type s_c strongly believes $R_{-c}^1(\mathcal{T})$. So, $Q_c^2 \subseteq \text{proj}_{S_c} R_c^2(\mathcal{T})$. ■

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

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

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

Observe that, by definition of an m-BRS, we can always find mappings $j_c^n : Q_c^n \to \mathcal{C}(Q_{-c}, \mathcal{E}_c)$ satisfying condition (j.b). But, condition (j.a) is stronger than that required by an m-BRS. If we find mappings $j_c = (j_c^1, \ldots, j_c^m)$ satisfying these requirements, we say that j_c **strongly justifies the m-BRS for player** c or j_a **and** j_b **strongly justify the m-BRS**. Theorem 6.1 follows from the following two propositions.

Proposition C.1. *Fix an m-BRS $(Q^0, \ldots, Q^m)$ satisfying the strong justification property. Then there exists a type structure $\mathcal{T}$ 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-BRS 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-BRS $(Q^0, \ldots, Q^m)$ satisfying the strong justification property. Thus, for each player c, there are mappings $j_c = (j_c^1, \ldots, j_c^m)$ that strongly justify the m-BRS.

Description of the Type Structure For each player c and each $n = 1, \ldots, m$, set $U_c^m \equiv Q_c^m$ and write $v_c^n : Q_c^n \to U_c^n$ for the identity map. The type set for player c will be $T_c = \bigsqcup_{n=1}^m U_c^n$. We will refer to types in U_c^n as the n-types for player c.

It will be convenient to specify the **diagonal** of $Q_c^n \times U_c^n$. This will be given by

$$\text{diag}_c^n = \bigcup_{s_c \in Q_c^n} ([s_c] \times \upsilon_c^n([s_c]))$$

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 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: 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 ($\leq n-1$) 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 $k = n-1, \ldots, m$ and some $(r_c, \upsilon_c^k(s_c)) \in \mathbb{BR}[j_c^k(s_c)] \times \{\upsilon_c^k(s_c)\} \subseteq \mathbb{Q}_c^k$. 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-1}.

(ii) If $k = n, \ldots, m$, then $\upsilon_c^k(s_c)$ strongly believes R_{-c}^{n-1}.

$\mathbf{n = 2}$: 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: $R_{-c}^1 = \bigcup_{k=1}^m \mathbb{Q}_{-c}^k$ and $Q_{-c}^1 = \text{proj}_{S_{-c}} \bigcup_{k=1}^m \mathbb{Q}_{-c}^k = \text{proj}_{S_{-c}} R_{-c}^1(\mathcal{T})$ (Lemma C.1).

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^1_{-c} is non-singleton. Observe that, in this case,

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

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, $v^1_c(s_c)$ 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,

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

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

$\mathbf{n \geq 3}$: Let $n = 3, \ldots, m$ and suppose the result was shown for $n - 1$. Fix some $k = n - 1, \ldots, m$ and some $(r_c, v^k_c(s_c)) \in \mathbb{BR}[j^k_c(s_c)] \times \{v^k_c(s_c)\} \subseteq \mathbb{Q}^k_c$. We show (i)-(ii).

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

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

Since $\bigcup_{l=n-1}^{k-1} \text{diag}^l_{-c} \subseteq \bigcup_{l=n-1}^{m} \mathbb{Q}^l_{-c}$ and, by the induction hypothesis, $R^{n-1}_{-c} = \bigcup_{l=n-1}^{m} \mathbb{Q}^l_{-c}$, it follows that $\beta_c(v^k_c(s_c))(R^{n-1}_{-c}|S_{-c}(h) \times T_{-c}) = 1$, 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.[20]

Observe that $r_c^* \notin [s_c]$. Also observe that r_c^* is optimal 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 optimal under $p_c(\cdot|S_{-c}(h^*))$ given $S_c(h^*)$, it follows that r_c^* is also optimal 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 optimal 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 optimal under $p_c(\cdot|S_{-c}(h))$ given $S_c(h)$, it follows that r_c^* is also optimal 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 optimal under $p_c(\cdot|S_{-c}(h^*))$ given $S_c(h^*)$, Definition 2.2 and Lemma A.3 give that r_c^* is optimal 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 optimal 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,

[20] 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.)

for any s_{-c} with $p_c(s_{-c}|S_{-c}(h^*)) > 0$, $\zeta(r_c^*, s_{-c}) = \zeta(r_c, s_{-c})$.

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 μ_c 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 optimal 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 optimal under $p_c(\cdot|S_{-c}(h))$ given $S_c(h)$ that r_c is also optimal 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 $\zeta(s_c^*, s_{-c}^*) = \zeta(r_c, s_{-c}^*)$ implies $r_c \in S_c(h^*)$.*

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 $s_c \in S_c \backslash [s_c^*]$.

Since $S_{-c} \backslash [s_{-c}^*] \neq \emptyset$, we can choose $r_{-c}^* \in S_{-c} \backslash [s_{-c}^*]$. 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 q_c^ε 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 q_c^ε that satisfies these conditions does not strongly believe $[s_{-c}^*]$. We show that we can choose $\varepsilon > 0$ so that $\mathbb{BR}[\mathrm{q}_c^\varepsilon] = [s_c^*]$.

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 optimal 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 $S(h^*) \subsetneq S(h)$. In that case, $q_c^\varepsilon(\cdot|S_{-c}(h)) = q_c^\varepsilon(\cdot|S_{-c})$. Second, suppose that $S(h) \subseteq S(h^*)$. 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, $s_c^* \in S_c(h^*) \cap S_c(h)$ 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 NRC. It shows Proposition 7.1. As a corollary, any game satisfying NRC is generic. The second half of the appendix shows that PI games satisfying NRT are generic.

D.1 No Relevant Convexities

The proof procedes as follows: Given a strategy s_c^* and an array $\mathrm{p}_c \in \mathcal{A}(S_{-c}, \mathcal{E}_c)$ with $s_c^* \in \mathbb{BR}[\mathrm{p}_c]$, we can construct a canonical CPS. Under that CPS, s_c^* remains a sequential best response. Moreover, the CPS preserves strong belief. We then show that, if the game satisfies NRC, then we can choose the CPS so that the set of best responses is simply $[s_c^*]$. This completes the proof of Proposition 7.1.

Lemma D.1. *Fix a strategy* s_c^* *and some array* $\mathrm{p}_c \in \mathcal{A}(S_{-c}, \mathcal{E}_c)$ *with* $s_c^* \in \mathbb{BR}[\mathrm{p}_c]$*. Then there exists a* ***canonical CPS for*** (s_c^*, p_c)*, viz.* $\mathrm{q}_c \in \mathcal{C}(S_{-c}, \mathcal{E}_c)$*, so that the following hold:*

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

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

We inductively construct the **canonical CPS for** (s_c^*, p_c), viz. $\mathrm{q}_c = (q_c(\cdot|S_{-c}(h)) : h \in H_c \cup \{\phi\})$ as follows: Let $H_c^0 = H_c \cup \{\phi\}$. 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 $h \in H_c^{k+1}$ with $S_{-c}(h^{k+1}) = S_{-c}(h)$, 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 $h \in H_c^{k+1}$ 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

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

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.* (Note, they depend on both p_c and s_c^*.) 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.2. *If* $h \in \overline{H}_c^k$ *and* $s_c^* \in S_c(h)$*, then* $S(h) \subseteq S(h^k)$*.*

Proof. Fix $h \in \overline{H}_c^k$ with $s_c^* \in S_c(h)$. Then, by construction, $s_c^* \in S_c(h) \cap S_c(h^k) \neq \emptyset$. Suppose, contra hypothesis, that $S(h)$ is not contained in $S(h^k)$. By perfect recall, either $S(h^k) \subsetneq S(h)$ or $S(h) \cap S(h^k) = \emptyset$. First, assume that $S(h^k) \subsetneq S(h)$. Again employing perfect recall,

$$S(h^k) = S_c(h^k) \times S_{-c}(h^k) \subsetneq S_c(h) \times S_{-c}(h) = S(h).$$

Using the fact that $S_{-c}(h^k) \subseteq S_{-c}(h)$ and Property (ii) of the construction, $S_{-c}(h^k) = S_{-c}(h)$. So, $S_c(h^k) \subsetneq S_c(h)$. But then, by Property (iii) of the construction, $S_c(h) \cap S_c(h^k) = \emptyset$, a contradiction.

Second, assume that $S(h) \cap S(h^k) = \emptyset$. Since $h \in \overline{H}_c^k$, $\emptyset \neq S_{-c}(h) \subseteq S_{-c}(h^k)$. It follows from Lemma A.2 that $S_c(h) \cap S_c(h^k) = \emptyset$, a contradiction. ■

We prove Lemma D.1 by showing that (i) $[s_c^*] \subseteq \mathbb{BR}[\mathrm{q}_c]$, and (ii) if p_c strongly believes E_{-c}, then q_c strongly believes E_{-c}.

Lemma D.3. $[s_c^*] \subseteq \mathbb{BR}[\mathrm{q}_c]$.

Proof. 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.2-A.3 that s_c^* is optimal under $q_c(\cdot|S_{-c}(h))$ given all strategies in $S_c(h)$. ■

Lemma D.4. *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 $q(E_{-c}|S_{-c}(h)) = 1$. ■

Say (s_c^*, p_c) satisfies **Property [*]** if the following holds:

> **Property [*]**: For each $h \in H_c$ with $s_c^* \in S_c(h)$, if $r_c \in S_c(h)$ is optimal under $p_c(\cdot|S_{-c}(h))$ among strategies in $S_c(h)$, 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.5. *Fix a strategy s_c^* and some array $\mathrm{p}_c \in \mathcal{A}(S_{-c}, \mathcal{E}_c)$ with $s_c^* \in \mathbb{BR}[\mathrm{p}_c]$. Suppose (s_c^*, p_c) satisfies Property [*]. Then, the canonical CPS for (s_c^*, p_c), viz. $\mathrm{q}_c \in \mathcal{C}(S_{-c}, \mathcal{E}_c)$, satisfies:*

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

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

Proof. By Lemma D.1, it suffices to show that $\mathbb{BR}[\mathrm{q}_c] \subseteq [s_c^*]$. Fix some $r_c \in \mathbb{BR}[\mathrm{q}_c] \backslash [s_c^*]$. Then there is an information set $h \in H_c$ so that $s_c^*, r_c \in S_c(h)$ and $s_c^*(h) \neq r_c(h)$. Let k be such that $h \in \overline{H}_c^k$ and note that r_c is a optimal under $q_c(\cdot|S_{-c}(h^k)) = p_c(\cdot|S_{-c}(h^k))$ given $S_c(h^k)$. Fix some $s_{-c} \in S_{-c}(h) \subseteq S_{-c}(h^k)$ such that $q_c(s_{-c}|S_{-c}(h)) > 0$. Observe that $\zeta(s_c^*, s_{-c}) \neq \zeta(r_c, s_{-c})$ and $p_c(s_{-c}|S_{-c}(h^k)) > 0$. This contradicts the fact that (s_c^*, p_c) satisfies Property [*]. ■

Lemma D.6. *Fix a game that satisfies NRC. Let s_c^* and $\mathrm{p}_c \in \mathcal{A}(S_{-c}, \mathcal{E}_c)$ be such that $s_c^* \in \mathbb{BR}[\mathrm{p}_c]$. Then there exists an array $\hat{\mathrm{p}}_c \in \mathcal{A}(S_{-c}, \mathcal{E}_c)$ so that*

(i) $(s_c^*, \hat{\mathrm{p}}_c)$ *satisfies Property [*],*

(ii) $s_c^* \in \mathbb{BR}[\hat{\mathrm{p}}_c]$, *and*

(iii) p_c *strongly believes* E_{-c} *if and only if* $\hat{\mathrm{p}}_c$ *strongly believes* E_{-c}*.*

Proof. For each $h \in H_c$ with $s_c^* \in S_c(h)$, we can choose $\hat{p}_c(\cdot|S_{-c}(h))$ so that (a) $r_c \in S_c(h)$ is optimal under $\hat{p}_c(\cdot|S_{-c}(h)$ amongst all strategies in $S_c(h)$ if and only if r_c supports s_c^* given $(\mathrm{Supp}\, p_c(\cdot|S_{-c}(h)), h)$, and (b) $\mathrm{Supp}\, \hat{p}_c(\cdot|S_{-c}(h)) = \mathrm{Supp}\, p_c(\cdot|S_{-c}(h))$. (See Lemmata D.2-D.3-D.4 in Brandenburger, Friedenberg and Keisler, 2008.) Requirement (i) follows from the construction and NRC. Requirements (ii)-(iii) follow immediately from the construction. ■

Proof of Proposition 7.1. Immediate from Lemmata D.5-D.6. ■

Corollary D.1. *If a game satisfies NRC, then it is generic.*

Proof. Fix some $s_c \in \mathbb{BR}[\mathrm{p}_c]$ for some $\mathrm{p}_c \in \mathcal{C}(S_{-c}, \mathcal{E}_c)$. By Lemmata D.5-D.6, there exists some $\mathrm{q}_c \in \mathcal{C}(S_{-c}, \mathcal{E}_c)$ with $[s_c] = \mathbb{BR}[\mathrm{q}_c]$. ■

One implication of Proposition 7.1 is that, when NRC is satisfied, we can forgo using CPS's and focus on arrays. This would not be the case absent NRC. The central difficulty comes from condition (BRP.3) of the m-BRS. Specifically, begin with a decreasing sequence of product sets $(Q^0, \ldots, Q^{m-1}, Q^m)$. In addition, suppose that $s_c \in Q_c^m$ 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. The canonical CPS $\mathrm{q}_c \in \mathcal{C}(S_{-c}, \mathcal{E}_c)$ will satisfy conditions (BRP.1)-(BRP.2). But, condition (BRP.3) may fail. The next example makes this point.

Example D.1. Consider the game in Figure D.1, which fails NRC. Write $h \in H_a$ 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 $\mathbb{BR}_a[\mathrm{p}_a] = \{O, I\text{-}U\}$.

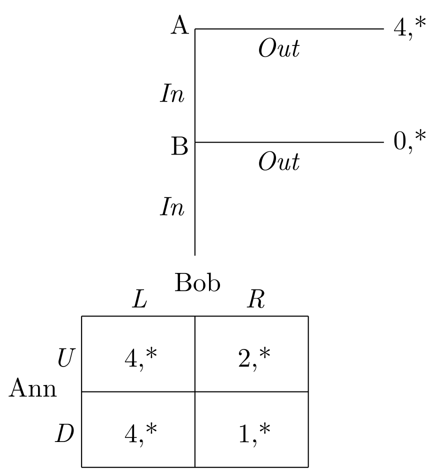

Figure D.1: Arrays Do Not Suffice

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

Definition D.1 (Battigalli, 1997)**.** A game satisfies **no relevant ties** (**NRT**) if $\pi_c(s_c, s_{-c}) = \pi_c(r_c, s_{-c})$ implies $\zeta(s_c, s_{-c}) = \zeta(r_c, s_{-c})$.

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 game that satisfies NRC satisfies NRT. But, a perfect information game can satisfy NRT, even if it fails NRC. Example D.3 below provides an example.

We first observe that a game can satisfy NRT, even though it is non-generic. We then give two classes of NRT games (one a subclass of the other) that are generic.

Example D.2. The game in Figure D.2 satisfies no relevant ties. Yet it is not generic: *Out* is a sequential best response 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 $\mathbb{BR}[\mathrm{q}_a] = [Out\,]$. □

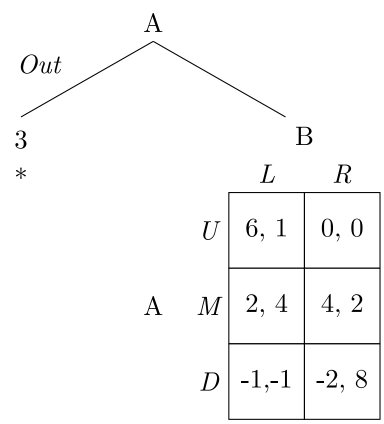

Figure D.2: No Relevant Ties

Note, in Example D.2, *Out* is justifiable, but not optimal under any CPS that involve degenerate beliefs—i.e., point beliefs. With this in mind:

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 $p(\omega|E) = 1$.

Note, in Definition D.2, ω may depend on E.

Definition D.3. Call a game **degenerately justifiable** if, whenever s_c is justifiable, there exists some degenerate CPS $\mathrm{p}_c \in \mathcal{C}(S_{-c}, \mathcal{E}_c)$ so that $s_c \in \mathbb{BR}[\mathrm{p}_c]$.

Example D.2 is not degenerately justifiable.

Proposition D.1. *A degenerately justifiable game that satisfies NRT is generic.*

Proof. Fix a degenerately justifiable game satisfying NRT and a justifiable strategy s_c. Then there exists a 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 $p_c(s_{-c}|S_{-c}(h)) = 1$. Since s_c is a sequential best response under p_c, $\pi_c(s_c, s_{-c}) \geq \pi_c(r_c, s_{-c})$. But, since $\zeta(s_c, s_{-c}) \neq \zeta(r_c, s_{-c})$, NRT implies $\pi_c(s_c, s_{-c}) > \pi_c(r_c, s_{-c})$. Thus, $r_c \notin \mathbb{BR}[\mathrm{p}_c]$. ■

Proposition D.2. *A perfect-information game satisfying no relevant ties is generic.*

The key will be the following:

Lemma D.7. *A perfect-information game satisfying NRT is degenerately justifiable.*

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.7. 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 $s_{-c}^h \in S_{-c}(h)$ so that $\pi_c(s_c, s_{-c}^h) \geq \pi_c(r_c, s_{-c}^h)$ for all $r_c \in S_c(h)$. Use the collection $(s_{-c}^h : h \in H_c \cup \{\phi\})$ to form a CPS p_c.

We will inductively define the measures $p_c(\cdot|S_{-c}(h))$. For each $S_{-c}(h)$ with $s_{-c}^{\{\phi\}} \in S_{-c}(h)$, set $p_c(s_{-c}^{\{\phi\}}|S_{-c}(h)) = 1$. 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 $p_c(s_{-c}^{h^*}|S_{-c}(h)) = 1$ for each $S_{-c}(h)$ with $s_{-c}^{h^*} \in S_{-c}(h)$. 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 a sequential best response 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 weakly) h so that $s_{-c} \in S_{-c}(h)$ and $p_c(s_{-c}^{h^*}|S_{-c}(h)) = 1$. Then, the claim follows from the fact that $S_c(h) \subseteq S_c(h^*)$ and the fact that $\pi_c(s_c, s_{-c}^h) \geq \pi_c(r_c, s_{-c}^h)$ for all $r_c \in S_c(h)$. ■

Proof of Proposition D.2. Immediate from Proposition D.1 and Lemma D.7. ■

Example D.3. The game in Figure D.3 is a perfect-information game satisfying NRT.[21] As such, it is generic. But, the conclusion of Proposition 7.1 does not hold. To see this, let (Q^0, Q^1, Q^2) be a decreasing sequence of product sets, so that

$$Q_a^1 \times Q_b^1 = \{O, L'R'', R'L'', R'R''\} \times \{o, lr', rr''\}$$

and

$$Q_a^2 \times Q_b^2 = \{O\} \times \{o\}.$$

Observe that Q^1 corresponds to the set of strategies that survive one round of EFR. Thus, (Q^0, Q^1) is a 1-BRS. We will argue that (Q^0, Q^1, Q^2) satisfies the requirements of Proposition 7.1, but is not a 2-BRS.

First observe that Ann's strategy O is a unique sequential best response under a CPS that *ex ante* assigns probability 1 to lr'; this CPS strongly believes Q_b^1. An array of Bob that strongly believes Q_a^1 must assign zero probability to $L'L''$ conditional upon Bob's first information set being reached. Thus, o is a sequential best response under an array p_b that strongly believes Q_a^1 if and only if, conditional upon Bob's first information set being reached, the array assigns probability $\frac{2}{3} : \frac{1}{3}$ to $L'R'' : R'L''$. Thus, (Q^0, Q^1, Q^2) satisfies the requirements of Proposition 7.1. But, it is not a 2-BRS: We can choose the array p_b to be a CPS. But, $\mathbb{BR}[\mathrm{p}_b] = \{o, lr', rr''\}$ is not contained in Q_b^2. Thus, (Q^0, Q^1, Q^2) fails the maximality criterion.

[21] A three-player version appears in Battigalli (1997).

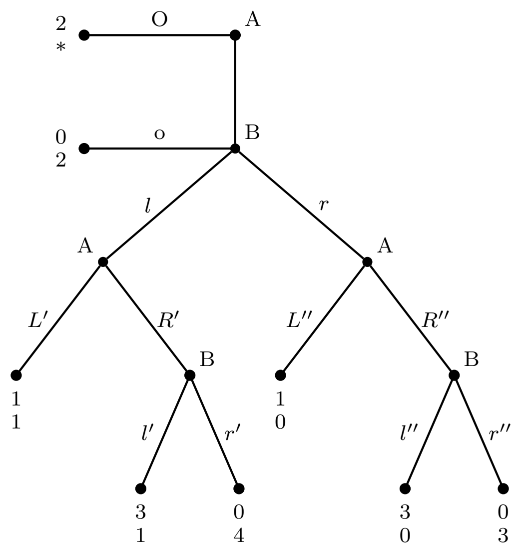

Figure D.3

Appendix E Proofs for Section 7

Proof of Proposition 7.2. Fix some game Γ. Let $\mathcal{Q} = (Q^0, Q^1, \ldots)$ be a BRS-sequence, i.e., for each finite m, $(Q^0, \ldots, Q^m)$ is an m-BRS. 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

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

For each BRS-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 $\overline{S}^{\overline{M}} \subseteq \overline{S}^{\infty}$. ■

Appendix F Centipede

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

Observation F.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 F.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 1 to $[\mathrm{out}, \mathrm{V}]_\ell$ at node $V-1$. (This uses Observation F.1, i.e., the fact that $[\mathrm{in}]_\ell \cap Q^1_\ell = \emptyset$.) Thus, $[\mathrm{in}]_{-\ell}$ is not a sequential 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 1 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 1 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 F.2. *Fix some* $m = 3, \ldots, V-1$. *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) $[\mathrm{out}, \mathrm{V}+2-\mathrm{m}]_\ell \cap Q^{m-1}_\ell = \emptyset$ *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. Note, since $V-1 \geq m \geq 3$, $V \geq 4$. So, by Lemma F.1, $[\mathrm{in}]_{-\ell} \cap Q^2_{-\ell} = \emptyset$. We repeatedly use this fact below.

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

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

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

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

So, any CPS p_ℓ that strongly believes $Q^2_{-\ell}$ 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 1 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 $[\mathrm{out}, V]_\ell \cap Q^3_\ell = \emptyset$.

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

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

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

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

So, any CPS $\text{p}_{-\ell}$ that strongly believes Q^3_ℓ must have

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

(That is, *ex ante*, $\text{p}_{-\ell}$ assigns probability 1 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 $p_{-\ell}([\text{out}, \tilde{v}+1]_{-\ell} | S_\ell(\tilde{v})) = 1$. Thus, at node $\tilde{v}$, $[\text{out}, \tilde{v}]_{-\ell}$ is a unique best response. So certainly $[\text{out}, V-2]_{-\ell} \cap Q^4_{-\ell} = \emptyset$. ■

Corollary F.1. *If $V = m$, then either $Q^V_a = [\text{out}, 1]_a$ or $Q^V_a = \emptyset$.*

Proof. We show the result for V odd. (The case of V even is analogous.) If $Q^{V-2}_a \notin \{[\text{out}, 1]_a, \emptyset\}$, the claim is immediate. So suppose otherwise. By Observation F.1, $[\text{in}]_a \notin Q^{V-2}_a$. By Lemma F.1, for each $m \leq V-2$ odd, $[\text{out}, V+3-m]_a \cap Q^{V-2}_a = \emptyset$. So,

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

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

Appendix G Section 9

Proof of Corollary 9.1. Certainly, (ii) implies (i). Suppose that (i) holds 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 6.1 and Observation 7.1 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

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

as desired. ■

Coarse Datasets McKelvey and Palfrey (1992) use the direct response method in their experimental design. In the main text, we suggested that doing so does not appear to be a significant hinderance to identification in their dataset. Here we expand on this point.[22]

McKelvey and Palfrey study 4 and 6 move Centipede games. Because there may be learning in the course of their experimental design, we focus on the last two periods of play. (Looking at earlier periods provides qualitatively similar conclusions.) We identify the first mover's rationality bound based on the minimum of the identified rationality bounds in the last two periods of play. When $V = 4$, all but one first-mover subject has a rationality bound identified as less than or equal to $3 = V - 1$. When $V = 6$, all first-mover subjects have a rationality bound identified as less than or equal to 5. In this latter case, 80% of first-mover subjects have a rationality bound identified as 3 or less.

References

Agranov, Marina, Elizabeth Potamites, Andrew Schotter and Chloe Tergiman. 2012. "Beliefs and endogenous cognitive levels: An experimental study." *Games and Economic Behavior* 75(2):449–463.

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

Alaoui, Larbi and Antonio Penta. 2017. "Reasoning about others' reasoning." `http://www.econ.upf.edu/~alaoui/alaoui_penta_reasoning_about_reasoning.pdf`.

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

Bagwell, Kyle and Garey Ramey. 1996. "Capacity, entry, and forward induction." *The Rand Journal of Economics* pp. 660–680.

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://u.arizona.edu/~afriedenberg/firr-online.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.

[22] An important caveat is in order. Recent work by Healy (2017) questions whether McKelvey and Palfrey (1992) did indeed implement a Centipede experiment.

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

Bradley, Stephen, Arnoldo Hax and Thomas Magnanti. 1977. *Applied mathematical programming*. Addison Wesley.

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.

Brandts, J., A. Cabrales and G. Charness. 2007. "Forward Induction and Entry Deterrence: An Experiment." *Economic Theory* 33(1):183–209.

Brandts, J. and C. Holt. 1995. "Limitations of Dominance and Forward Induction: Experimental Evidence." *Economics Letters* 49(4):391–395.

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

Cooper, D, Enrique Fatas, Antonio J Morales and Shi Qi. 2016. "Consistent depth of reasoning in level-k models." http://fae.ua.es/FAEX/wp-content/uploads/2018/02/Consistent_LevelK.pdf.

Cooper, R., D. DeJong, R. Forsythe and T. Ross. 1993. "Forward Induction in the Battle-of-the-Sexes Games." *American Economic Review* 83(5):1303–1316.

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 and Marciano Siniscalchi. 2014. "Epistemic game theory." *Handbook of Game Theory* 4.

Dufwenberg, Martin and Matt Van Essen. 2018. "King of the Hill: Giving backward induction its best shot." *Games and Economic Behavior* 112:125–138.

Friedenberg, A., W. Kets and T. Kneeland. 2016. "Is Bounded Rationality Driven by Limited Ability?". http://u.arizona.edu/~afriedenberg/car.pdf.

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

Healy, Paul J. 2017. "Epistemic Experiments: Utilities, Beliefs, and Irrational Play." https://www.asc.ohio-state.edu/economics/healy/papers/Healy-EpistemicExperiments.pdf.

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

Huck, Steffen and Wieland Müller. 2005. "Burning money and (pseudo) first-mover advantages: an experimental study on forward induction." *Games and Economic Behavior* 51(1):109–127.

Kets, W. 2010. "Bounded Reasoning and Higher-Order Uncertainty." https://wkets.org.

Kinderman, P., R. Dunbar and R. Bentall. 1998. "Theory-of-Mind Deficits and Causal Attributions." *British Journal of Psychology* 89(2):191–204.

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.

Mantovani, Marco. 2014. "Limited backward induction: foresight and behavior in sequential games.".

McKelvey, R. and T. Palfrey. 1992. "An Experimental Study of the Centipede Game." *Econometrica* 60(4):803–836.

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.

Reny, Philip J. 1993. "Common belief and the theory of games with perfect information." *Journal of Economic Theory* 59(2):257–274.

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." http://faculty.wcas.northwestern.edu/~msi661/.

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

Stiller, James and Robin IM Dunbar. 2007. "Perspective-taking and memory capacity predict social network size." *Social Networks* 29(1):93–104.