Theoretical Economics 16 (2021), 1605–1654 1555-7561/20211605

The implications of finite-order reasoning

Adam Brandenburger
Stern School of Business, New York University, and NYU Shanghai

Alexander Danieli
W.P. Carey School of Business, Arizona State University

Amanda Friedenberg
Department of Economics, University of Arizona

The epistemic conditions of rationality and mth-order strong belief of rationality (RmSBR; Battigalli and Siniscalchi (2002)) formalize the idea that players engage in contextualized forward-induction reasoning. This paper characterizes the behavior consistent with RmSBR across all type structures. In particular, in a class of generic games, R$(m-1)$SBR is characterized by a new solution concept we call an m-best response sequence (m-BRS). Such sequences are an iterative version of extensive-form best response sets (Battigalli and Friedenberg (2012)). The strategies that survive m rounds of extensive-form rationalizability are consistent with an m-BRS, but there are m-BRS's that are disjoint from the former set. As such, there is behavior that is consistent with R$(m-1)$SBR but inconsistent with m rounds of extensive-form rationalizability. We use our characterization to draw implications for the interpretation of experimental data. Specifically, we show that the implications are nontrivial in the three-repeated Prisoner's Dilemma and Centipede games.

Keywords. Epistemic game theory, strategic uncertainty, bounded reasoning, identifying reasoning.

JEL classification. D01, D03, D83.

Adam Brandenburger: adam.brandenburger@stern.nyu.edu
Alexander Danieli: Alexander.Danieli@asu.edu
Amanda Friedenberg: amanda@amandafriedenberg.org
We thank 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 Grant SES-1358008 for financial support. Amanda Friedenberg thanks the UCL Economics Department for its unbounded hospitality. Previous versions of this paper were circulated under the titles: "How Many Levels Do Players Reason: An Observational Challenge and Solution" and "Identification of Reasoning about Rationality."

Available at https://econtheory.org. https://doi.org/10.3982/TE2889

1. Introduction

Suppose that each player is rational, each player thinks other players are rational, and so on ad infinitum. What are the implications for behavior? This fundamental question has organized the epistemic game theory literature. It has been asked in the context of both strategic-form and extensive-form games. It has been asked when rationality corresponds to ordinary subjective expected utility maximization, when rationality incorporates an admissibility requirement, and even when rationality departs from subjective expected utility maximization. Likewise, it has been asked when the word "thinks" corresponds to knowledge, to belief, and to many other modalities (e.g., assumption, initial belief, strong belief, etc).

This paper focuses on a bounded version of the basic question, "What are the implications for behavior if each player is rational, each player thinks other players are rational, and so on up to m levels, but not $(m+1)$ levels?" We address this question in the context of extensive-form games, where players engage in "contextualized forward-induction reasoning." We go on to spell out implications for the interpretation of experimental data.

Consider the game in Figure 1, Battle of the Sexes with an Outside Option (BoSOO).[1] A natural first approach would involve some notion of iterated dominance—here, extensive-form rationalizability (EFR; Pearce (1984)). EFR gives the standard forward-induction prediction: Ann plays *In-U* and Bob plays *L*. In particular, the strategy *In-D* is dominated, but *In-U* is not. So, under EFR, Bob believes Ann will play *U* conditional on Battle of the Sexes (BoS); therefore, he plays *L*. Ann expects this and so plays *In-U*. With this, the analyst may be tempted to conclude that, if Ann plays *Out*, it is not the case that Ann is rational and reasons (in the sense of forward induction) two levels about Bob's rationality. However, we will see that this conclusion is incorrect.

To see this, we adopt the epistemic approach. In particular, we expand the description of the game to include the players' hierarchies of beliefs about the play of the game.

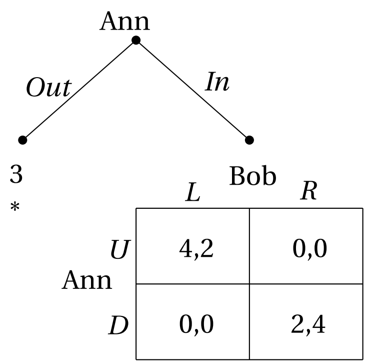

Figure 1. Battle of the Sexes with an Outside Option.

[1]Here, and in other examples, the asterisk (∗) indicates that the payoff is irrelevant.

To understand why this is needed, note that we will be interested in which of Ann's strategies are rational. The strategy *Out* is rational—i.e., a best response—if Ann believes that Bob will play *R*; it is irrational if she believes that Bob will play *L*. Thus, we cannot specify whether *Out* is rational or irrational without specifying Ann's beliefs about Bob's play of the game.

In fact, we need to specify hierarchies of conditional beliefs about the play of the game. To see this, consider the case where Ann believes that Bob plays *R*. This strategy is rational for Bob if, conditional on BoS being played, Bob believes that Ann plays *D*. But, this same strategy is irrational for Bob if, conditional on BoS being played, Bob believes Ann plays *U*. Thus, to specify whether Ann is rational and believes that Bob is rational, we need to specify both (i) Ann's belief about the strategy Bob plays, and (ii) her belief about Bob's conditional belief about her own play (where the conditioning is on BoS). Continuing along these lines, we need to specify Ann's hierarchies of conditional beliefs about the play of the game.

To do this, we build on the framework and analyses in Battigalli and Siniscalchi (2002) and Battigalli and Friedenberg (2012). We describe the situation by what is called an epistemic game. This is the game itself (e.g., BoSOO) plus an epistemic type structure. Each type in an epistemic type structure is associated with a conditional probability system, which describes a player's belief at every information set—including beliefs at information sets that are, ex ante, not expected to be reached. For instance, in BoSOO, if Bob believes Ann plays *Out*, he does not expect his information set to be reached. Nevertheless, a conditional probability system specifies his belief conditional on this event. This framework allows us to specify the requirement that each player is rational, thinks the other is rational, etc.

Our baseline epistemic condition is rationality and common strong belief of rationality (RCSBR). There are two ingredients to this condition: (extensive-form) rationality and strong belief. Extensive-form rationality requires that a player chooses a strategy that, at each information set, maximizes her subjective expected utility given her type's conditional probability system. For strong belief, note that, in an extensive form, we cannot simply require that a player "believes" an event: We must specify where, in the game, the player believes that event.[2] Strong belief requires that a player begins the game with a belief that the event is true and maintains that hypothesis as long as it is not contradicted by evidence. Returning to BoSOO, if Bob strongly believes Ann is rational and reaching the BoS subgame is consistent with Ann's rationality, then he assigns probability 1 to her rationality conditional upon the BoS subgame. (If reaching the BoS subgame is inconsistent with Ann's rationality, strong belief of rationality requires only that Bob begin the game by assigning probability 1 to Ann's rationality.) Thus, under RCSBR, a player rationalizes past behavior as long as possible.

Battigalli and Siniscalchi (2002) and Battigalli and Friedenberg (2012) both study the behavioral implications of RCSBR. However, they do so under importantly different assumptions. Battigalli and Siniscalchi (2002) focus on the case of a complete type structure (Brandenburger (2003)). Loosely, this corresponds to an assumption that players

[2]For this reason, we follow the modern literature and refrain from using the phrase "mutual belief of rationality" for the idea that a player thinks or reasons that the other player is rational.

have all possible hierarchies of beliefs. (See Friedenberg (2010).) By contrast, Battigalli and Friedenberg (2012) depart from the complete type structure assumption. The idea is that there may be a context to the strategic situation and that context may frame the hierarchies of beliefs that the players consider possible. For instance, history or social/cultural norms may impact the players' beliefs. (See, the discussion in Brandenburger et al. (2008, Section 2.8).) Under this perspective, a complete type structure represents a special "context-free" case in which there are no restrictions on the players' beliefs.

The distinction between the complete (context-free) versus incomplete (contextualized) cases has important implications for RCSBR (and, in particular, for RCSBR behavior). Battigalli and Siniscalchi (2002) show that, for a complete type structure, RCSBR behavior corresponds to EFR. So, in BoSOO, Ann plays *In-U* and Bob plays *L*. By contrast, Battigalli and Friedenberg (2012) characterize RCSBR across all type structures. In particular, it is characterized by a solution concept they call an extensive-form best response set (EFBRS), a natural extensive-form analog to Pearce's (1984) (strategic-form) best response solution concept. The EFR strategy set is one EFBRS, but there are others. For instance, in BoSOO, there is an EFBRS where Ann plays *Out*. This can arise if, for instance, history leads Ann to maintain the hypothesis (throughout the game) that Bob is "tough," and so "goes for his best outcome" and plays *R*. (This is formalized as an incomplete type structure. See Example 3.1.) This would lead any rational Ann to play *Out*. Moreover, this is consistent with contextualized forward-induction reasoning. In particular, when Bob has such a reputation, the observation of *In* is inconsistent with Ann being rational. As such, Bob can both strongly believe Ann is rational and, conditional on BoS, believe that playing *R* would lead to his best outcome.

This paper studies the context-dependent behavioral implications of rationality and mth-order strong belief of rationality, namely RmSBR. (We take R0SBR to mean just rationality.) As in Battigalli and Friedenberg (2012), we seek to characterize the RmSBR behavior across all type structures. This characterization is of particular interest when the analyst does not know the relevant type structure. As such, to infer that a strategy is consistent versus inconsistent with RmSBR, the analyst needs the RmSBR prediction across all type structures.

There is a natural approach to characterize the RmSBR behavior across all type structures: "unwrap" the EFBRS solution concept, thereby converting it from a fixed-point definition into an iterative definition. This leads to a sequence of predictions, viz. $(Q^0, \dots, Q^m)$. We call such a sequence an m-(extensive-form) best response sequence (m-BRS). For a given type structure, the sequence of behavior consistent with R0SBR, …, R$(m-1)$SBR gives rise to an m-BRS. (See Proposition 6.1.) However, there are m-BRSs that, in a certain sense, do not reflect R$(m-1)$SBR behavior. (Examples 6.1 and 6.2 make this precise.) Nonetheless, Theorem 6.1 establishes that in a class of generic games, the m-BRSs characterize R0SBR, …, R$(m-1)$SBR behavior. (See Definition 6.4 for the meaning of "generic.")

There is behavior consistent with an m-BRS but inconsistent with m-EFR, i.e., the strategies that survive m rounds of EFR. More precisely, we can have an m-BRS $(Q^0, \dots, Q^m)$ with Q^m disjoint from the set of m-EFR strategies. This is to be expected given that an EFBRS may be disjoint from EFR. Both scenarios can arise because strong belief is non-monotonic.[3] (The discussion following Examples 4.1 and 4.2 elaborates on this point.) This said, the m-EFR strategies are consistent with one m-BRS.

What does this mean for identifying levels of forward-induction reasoning from experimental data? In BoSOO, there is a 3-BRS in which Ann plays *Out*—something that is inconsistent with 3-EFR. However, this fact is not surprising, given that *Out* is also consistent with an EFBRS. In fact, in BoSOO, all undominated strategies are consistent with an EFBRS. But there are other games where the distinction between m-BRSs and m-EFR is both informative and important for the interpretation of experimental data. We give two such examples.

First, we look at the three-repeated Prisoner's Dilemma. Round for round, EFR gives the expected theoretical prediction. Round 1 has players defect at all third-period histories, round 2 has players defect at all second-period histories, and round 3 has players defect initially. The m-BRSs give the same path of play, but they allow for additional strategies. As such, the inference from observed behavior is more subtle. In particular, if the experimentalist uses the direct response method and observes a subject choose to cooperate in the second period, then—*irrespective of the behavior of the other subjects*—she cannot conclude that the behavior is inconsistent with rationality and strong belief of rationality. If, instead, the experimentalist uses the strategy method and observes a subject choose to cooperate at every second-period history, then she can conclude that the behavior is inconsistent with rationality and strong belief of rationality.

Second, we turn to the Centipede game (Rosenthal (1981)), where EFR corresponds round-for-round to backward induction. Thus, one might conjecture that, if we observe a subject play *in*, the behavior indicates a bound on strategic reasoning, and the longer the subject plays *in*, the lower that bound. This intuition is incomplete, at least if we take "strategic reasoning" to reflect RCSBR. In particular, we will see that the intuition is correct for the first mover, but incorrect for the second mover.

The paper proceeds as follows. Sections 2–4 introduce the formalism and epistemic conditions. Section 5 reviews the characterization of RCSBR. It is used as a benchmark result to motivate our main result. Section 6 defines the m-BRS concept and shows the main result: a characterization of R$(m-1)$SBR sequences in terms of m-BRSs. Section 7 uses the m-BRS concept to analyze the three-repeated Prisoner's Dilemma and Centipede. Section 8 concludes with a discussion. It highlights several issues that are important for applying the m-BRS concept, including which games are generic, termination of the m-BRS procedure, and when arrays can replace conditional probability systems. Proofs can be found in the Appendices.

2. Extensive-form games

Write Γ for a finite two-player extensive-form game with perfect recall, with (potentially) simultaneous moves, and without moves by nature. The players are a (Ann) and

[3]Note that this is different from (strategic-form) rationalizability: Best response sets are contained in the rationalizable strategy set. This follows from the fact that belief is monotonic, i.e., if E is believed and $E \subseteq F$, then F is believed.

b (Bob).[4] Write c for an arbitrary player in $\{a, b\}$ and write $-c$ for the player in $\{a, b\}\backslash\{c\}$. The underlying game tree has a set of nonterminal nodes (or vertices) V and a set of terminal nodes Z. Write $\phi \in V$ for the initial node. As in Osborne and Rubinstein (1994), we often identify nodes with histories. Let H_c be the set of information sets of c. The set of information sets is $H = H_a \cup H_b$. 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 $S = S_a \times S_b$. Assume the game is *nontrivial* in the sense that $|S_a|, |S_b| \geq 2$. 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). Say (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 $S_c(h) = \mathrm{proj}_{S_c} S(h)$. If a strategy $s_c \in S_c(h)$, then we say that s_c *allows* $h \in H$.

Player c's strategic-form payoff function is given by $\pi_c = \Pi_c \cdot \zeta$. We extend π_c to $S_c \times \mathcal{P}(S_{-c})$ in the usual way, i.e., $\pi_c(s_c, \mu) = \sum_{s_{-c}} \pi_c(s_c, s_{-c})\mu(s_{-c})$. Say $s_c \in X_c$ is *optimal under* $\mu \in \mathcal{P}(S_{-c})$ *given* X_c if $\pi_c(s_c, \mu) \geq \pi_c(r_c, \mu)$ for each $r_c \in X_c$.

3. Type structures

This section uses type structures to implicitly model the players' hierarchies of beliefs about the play of the game. In defining such structures, we use the language of conditional probability systems. To understand the importance of doing so, refer back to BoSOO. Suppose Bob begins the game by assigning probability 1 to Ann playing *Out*. Then, conditional upon observing Ann play *In*, Bob can no longer hold that belief. So, to specify which strategy is rational for Bob, we must specify his belief conditional upon Ann playing *In*.

Conditional probability systems Fix a compact metric space 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 $(\Omega, \mathcal{E})$ a *conditional probability space* if $\mathcal{E} \subseteq 2^{\Omega}\backslash\{\emptyset\}$ is a finite set, where each $E \in \mathcal{E}$ is Borel. The collection $\mathcal{E}$ is a set of *conditioning events*.

Definition 3.1. 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 $p(E|E) = 1$.

Definition 3.2. 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

[4]The analysis extends to three or more players, provided two assumptions hold. First, a player may have a correlated belief about other players. Second, players engage in so-called correlated rationalization, e.g., if they observe irrational behavior by one player, they are willing to entertain a hypothesis that other players may also be irrational.

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 CPSs 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.

Type structures In our analysis, player c's set of conditioning events corresponds 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, namely, $S_b(\phi) = S_b$. She also has conditioning events $S_b(h)$ corresponding to each information set $h \in H_a$ at which she moves.

Definition 3.3. A Γ-*based type structure* is some $\mathcal{T} = (\Gamma; T_a, T_b; \beta_a, \beta_b)$, where

(i) T_c is a compact metric *type space* for player c, and

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

So each type of Ann, viz. t_a, is associated with a CPS $\beta_a(t_a)$ on $(S_b \times T_b, \mathcal{E}_a \otimes T_b)$, and similarly for Bob. When Γ is a simultaneous-move game, the set of CPSs is the set of probability measures and so $\beta_c : T_c \to \mathcal{P}(S_{-c} \times T_{-c})$.

For any given game Γ, there are infinitely many Γ-based type structures. Write $\mathbb{T}(\Gamma)$ for the family of Γ-based type structures. Battigalli and Siniscalchi (1999) construct a canonical type structure that 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 (what is called) common full belief. (See Appendix A in Battigalli and Friedenberg (2009) for a formal treatment.) The following example informally illustrates such an assumption.

Example 3.1. Consider BoSOO. Suppose it is commonly understood that "Bob is tough" and, so, whenever a BoS game is played, he attempts to go for his best option and play R. In particular, the following conditions hold:

Tough 1. At both the start of the game and conditional on BoS, Ann believes that Bob plays R.

Tough 2. At both the start of the game and conditional on BoS, Bob believes Tough 1.

Tough 3. At both the start of the game and conditional on BoS, Ann believes Tough 2.

And so on. This is a restriction on the hierarchies of beliefs that the players consider possible. There are no additional restrictions on the players' beliefs.

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 first requirement says that, at the start of the game, each type of Ann assigns probability 1 to "Bob plays R." By the conditioning requirement, this implies that, conditional on BoS, each type continues to assign probability 1 to "Bob plays R." The second requirement says that, for each CPS that satisfies the first requirement, there is a type of Ann that holds that belief. Likewise, the third requirement says that for each CPS of Bob, there is a type of Bob that holds that belief. The second and third requirements capture the idea that there are no additional restrictions on the players' beliefs. The fact that such a type structure exists follows from Battigalli and Friedenberg (2009). ◊

Epistemic game For a given game Γ, write $\mathbb{T}(\Gamma)$ for the family of Γ-based type structures. Since Γ is nontrivial, 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.

In what follows, we fix an extensive-form game Γ. So, each epistemic game can be identified with a type structure in $\mathbb{T}(\Gamma)$. As such, we often conflate "type structure" with "epistemic game." No confusion should result.

4. Epistemic conditions

Fix an epistemic game $(\Gamma, \mathcal{T})$. It induces a set of *states* $S_a \times T_a \times S_b \times T_b$. A state describes the strategies played and the beliefs held. We focus on the set of states that satisfy rationality and mth-order strong belief of rationality. We begin with rationality.

Definition 4.1. 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 $s_c \in \mathbb{BR}[\mathrm{p}_c]$ if and only if s_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 $\mathrm{marg}_{S_{-c}} \beta_c(t_{-c})$ 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 4.2. 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 $\mathrm{marg}_{S_{-c}} \beta_c(t_c)$.

Definition 4.3 (Battigalli and Siniscalchi (2002)). Say an array $\mathrm{p} \in \mathcal{A}(\Omega, \mathcal{E})$ *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 4.4. 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 (relative to the type structure $\mathcal{T}$). See Battigalli and Siniscalchi (2002) and Battigalli and Friedenberg (2012) for a more complete discussion.

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). Inductively define sets $R_a^m(\mathcal{T})$ and $R_b^m(\mathcal{T})$ by

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

Set $R_c^\infty(\mathcal{T}) = \bigcap_{m\geq 0} R_c^m(\mathcal{T})$. Write $R^m(\mathcal{T}) = R_a^m(\mathcal{T}) \times R_b^m(\mathcal{T})$ and $R^\infty(\mathcal{T}) = R_a^\infty(\mathcal{T}) \times R_b^\infty(\mathcal{T})$.

Definition 4.5. The set $R^{m+1}(\mathcal{T})$ is the set of strategy-type pairs (in $(\Gamma, \mathcal{T})$) at which there is rationality and *mth-order strong belief of rationality* (*RmSBR*). The set $R^\infty(\mathcal{T})$ is the set of strategy-type pairs (in $(\Gamma, \mathcal{T})$) at which there is *rationality and common strong belief of rationality* (*RCSBR*).

Note that the set of rational strategy-type pairs depends on the epistemic game $(\Gamma, \mathcal{T})$. In this paper, we fix Γ but not necessarily the associated type structure. As such, we write $R^1(\mathcal{T})$ to indicate the set of states at which there is rationality in the epistemic game associated with $\mathcal{T}$, Analogously for $R^m(\mathcal{T})$.

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 the set of RmSBR predictions is the set of strategies that survive $(m+1)$ rounds of extensive-form rationalizability (EFR; Pearce (1984)). EFR sequentially eliminates strategies that are not sequential best responses. Battigalli and Siniscalchi (2002, Proposition 6) show that when the type structure $\mathcal{T}^*$ 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 set of RmSBR predictions may be disjoint from the set of strategies that survive $(m+1)$ rounds of EFR. The following examples illustrate this point.

Example 4.1. Consider BoSOO. There is no CPS so that *In-D* is a sequential best response. However, for each strategy $s_a \in \{Out, In\text{-}U\}$ (resp. $s_b \in \{L, R\}$), there is some CPS under which s_a (resp. s_b) is a 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 EFR_a^1 must assign probability 1 to *In-U* conditional on BoS; L is the unique sequential best response under any such CPS. As such,

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

So, *In-U* is the unique sequential best response under any CPS on S_b that strongly believes EFR_b^2. Thus,

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

As such, there is one EFR strategy profile, $(In\text{-}U, L)$.

Battigalli and Siniscalchi (2002) show that, if $\mathcal{T}^*$ is type-complete, then EFR corresponds round-for-round to RmSBR in the associated epistemic game. That is, for each m, $\mathrm{proj}_{S_a \times S_b} R^m(\mathcal{T}^*) = \mathrm{EFR}_a^m \times \mathrm{EFR}_b^m$. So the EFR predictions are the RmSBR predictions when the type structure is type-complete. ◊

Example 4.2. Again, consider BoSOO. Let $\mathcal{T}$ be the type structure from Example 3.1, representing the case where it is commonly understood that "Bob is tough." Now, for each $m \geq 1$, $\mathrm{proj}_{S_a \times S_b} R^m(\mathcal{T}) = \{Out\} \times \{L, R\}$. So, for each $m \geq 1$, there are types t_a^m and t_b^m so that $(Out, t_a^m, R, t_b^m) \in R^m(\mathcal{T})$. ◊

Let us review the difference between Examples 4.1 and 4.2. Begin with Example 4.2. 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 on 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" in the epistemic game $(\Gamma, \mathcal{T})$, viz. $R_a^1(\mathcal{T})$: At the initial node, the type assigns probability 1 to the event $R_a^1(\mathcal{T})$ and the event $R_a^1(\mathcal{T})$ is inconsistent with Ann playing *In*. Now turn to Example 4.1 and an associated complete type structure $\mathcal{T}^*$. There is also a type $t_b^{2*} \in T_b^*$ that, at the initial node, assigns probability 1 to $\{Out\} \times T_a^*$ and, conditional on Ann playing *In*, assigns probability 1 to $\{In\text{-}D\} \times T_a^*$. However, this type does not strongly believe the event that "Ann is rational" in the epistemic game $(\Gamma, \mathcal{T}^*)$, viz. $R_a^1(\mathcal{T}^*)$, since

$$R_a^1(\mathcal{T}^*) \cap \left(\{In\text{-}U, In\text{-}D\} \times T_a^*\right) \neq \emptyset$$

and, in particular, is contained in $\{In\text{-}U\} \times T_a^*$. The key is

$$\{Out\} = \mathrm{proj}_{S_a} R_a^1(\mathcal{T}) \subsetneq \mathrm{proj}_{S_a} R_a^1(\mathcal{T}^*) = \{Out, In\text{-}U\}.$$

As such, we can have a CPS that assigns positive probability to $\{In\text{-}D\} \times T_a$ conditional on BoS and strongly believes $R_a^1(\mathcal{T})$, but there is no CPS that assigns positive probability to $\{In\text{-}D\} \times T_a^*$ conditional on BoS and strongly believes $R_a^1(\mathcal{T}^*)$. This is possible because of the non-monotonicity of strong belief.

5. The EFBRS benchmark

We focus on the case where the analyst does not know the players' type structure. Thus, we are interested in characterizing the RmSBR predictions across all type structures. We begin with the RCSBR benchmark, where the analogous characterization is known.

Fix a type structure $\mathcal{T}$. The set of RCSBR predictions for $\mathcal{T}$ is

$$\mathrm{proj}_{S_a \times S_b} R^\infty(\mathcal{T}) = \mathrm{proj}_{S_a} R_a^\infty(\mathcal{T}) \times \mathrm{proj}_{S_b} R_b^\infty(\mathcal{T}).$$

Fix a predicted strategy $s_a \in \mathrm{proj}_{S_a} R_a^\infty(\mathcal{T})$ 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., $\mathrm{p}_a = \mathrm{marg}_{S_b} \beta_a(t_a)$. Observe that (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 $\mathcal{T}$, namely $\mathrm{proj}_{S_b} R_b^\infty(\mathcal{T})$. (This follows from the fact that t_a strongly believes

$R_b^\infty(\mathcal{T})$.) Finally, if r_a is also a sequential best response under p_a, r_a must be contained in $\mathrm{proj}_{S_a} R_a^\infty(\mathcal{T})$, i.e., r_a must be one of Ann's RCSBR predictions for $\mathcal{T}$. (This follows from the 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 $\mathrm{p}_c \in \mathcal{C}(S_{-c}, \mathcal{E}_c)$ so that

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

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

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

An EFBRS is a subset of strategies $Q_a \times Q_b$ that satisfies a certain fixed-point requirement: For each $s_a \in Q_a$, there exists a CPS p_a defined only on S_b 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, but are properties defined on the game Γ itself.

Proposition 5.1 (Battigalli and Friedenberg (2012)).

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

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

Corollary 5.1. *For each game Γ,*

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

Corollary 5.1 says that the union over all EFBRSs corresponds to the union of the RCSBR predictions. So a strategy s_c is consistent with RCSBR if and only if $s_c \in \mathrm{proj}_{S_c} S^\infty$.

Example 5.1. Return to Figure 1. We will show that

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

First, each EFBRS $Q^a \times Q^b \subseteq S^\infty$. To see this, fix an EFBRS $Q^a \times Q^b$. Since, at the start of the game, *In-D* is dominated, by property (i) of Definition 5.1, $Q^a \subseteq \{Out, In\text{-}U\}$. So if $In\text{-}U \in Q_a$, then $Q_b = \{L\}$. (This follows from properties (i) and (ii) of Definition 5.1.) Thus, $Q^a \times Q^b \subseteq S^\infty$. Next, $\{Out\} \times \{L, R\}$ and $\{In\text{-}U\} \times \{L\}$ are both EFBRSs.[5] Take $Q_a \times Q_b = \{Out\} \times \{L, R\}$. *Out* is the unique sequential best response under a CPS that, at each information set, assigns probability 1 to $\{R\} \subseteq Q_b$; L (resp. R) is the unique sequential best response under a CPS that initially assigns probability 1 to $Q_a = \{Out\}$ and,

[5] So is $\{Out\} \times \{R\}$.

conditional on BoS, assigns probability 1 to $\{In\text{-}U\}$ (resp. $\{In\text{-}D\}$). And, analogously for $\{In\text{-}U\} \times \{L\}$.

Finally, note that S^∞ is not itself an EFBRS. Observe that it is not a product set. But, more importantly, R is not a best response given a CPS that strongly believes $\text{proj}_{S_a} S^\infty = \{Out, In\text{-}U\}$. ♢

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 analog, 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 $R_b^3(\mathcal{T})$; that is, t_a does not strongly believe the event of the same order. Instead, t_a strongly believes the lower-order events $R_b^0(\mathcal{T})$, $R_b^1(\mathcal{T})$, and $R_b^2(\mathcal{T})$. 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 $R_b^0(\mathcal{T}), \dots, R_b^{m-1}(\mathcal{T})$. We will build off this fact to go from the EFBRS concept to an iterative property. That property applies to a decreasing sequence of product sets.

Definition 6.1. Say $(Q^0, \dots, Q^m)$ 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 $n = 0, \dots, m-1$, $Q^{n+1} \subseteq Q^n$.

Definition 6.2. Say $X = X_a \times X_b$ satisfies the (*extensive-form*) *best response property* relative to $(Q^0, \dots, Q^m)$ if $(Q^0, \dots, Q^m, X)$ 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 $Q_{-c}^0, \dots, Q_{-c}^m$, 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 $Q_{-c}^0, \dots, Q_{-c}^m$. (Note that $X_{-c} \subseteq Q_{-c}^m \subseteq Q_{-c}^{m-1} \subseteq \cdots \subseteq Q_{-c}^0$.)

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

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

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 nonempty $E_c \subseteq \mathcal{C}(S_{-c}, \mathcal{E}_c)$. An $(m+1)$-BRS is some $(Q^0, \dots, Q^m, Q^{m+1})$, where $(Q^0, \dots, Q^m)$ is an m-BRS and Q^{m+1} satisfies the best response property relative to $(Q^0, \dots, Q^m)$.[6] Thus, it is an iterative procedure that is a natural analog to the EFBRS. In fact, the following holds:

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

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

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

That is, the union over the m-BRSs provides an upper bound on the behavior consistent with R$(m-1)$SBR across all type structures. We denote the union over m-BRSs by S^m.

A natural analog to Corollary 5.1 is that

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

i.e., the union over the m-BRSs is the set of strategies consistent with R$(m-1)$SBR in some type structure. In fact, there is a natural conjecture that would imply Equation (1): For each m-BRS $(Q^0, \dots, Q^m)$, there exists a type structure $\mathcal{T}$ so that

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

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

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

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

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

Note this is a 2-BRS.[7] But we show that there is no type structure $\mathcal{T}$ with $Q^1 \subseteq \mathrm{proj}_S R^1(\mathcal{T})$ and $Q^2 = \mathrm{proj}_S R^2(\mathcal{T})$.

[6] Note that if $(Q^0, \dots, Q^m, Q^{m+1})$ is an $(m+1)$-BRS and $Q^m = Q_a^m \times Q_b^m$ with $Q_a^m = \emptyset$, then $Q_b^{m+1} = \emptyset$.

[7] 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

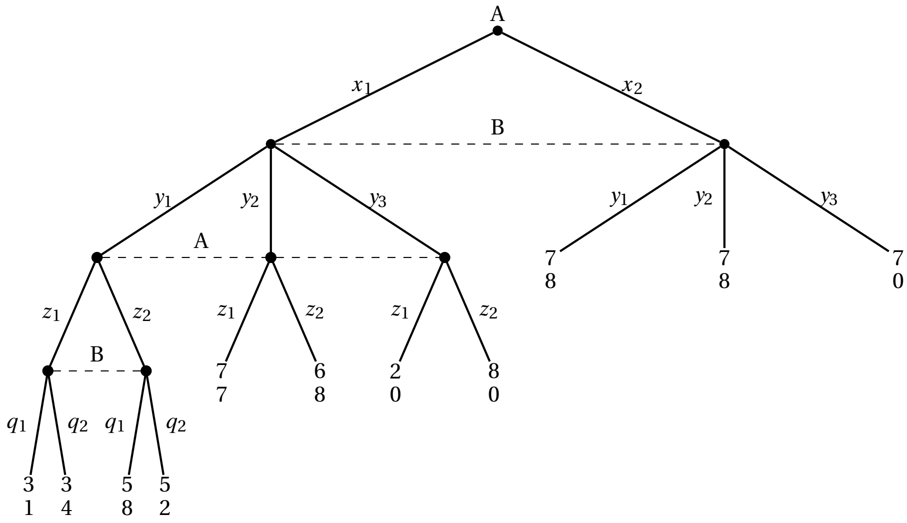

FIGURE 2. A 2-BRS that is not a R1SBR prediction.

Suppose otherwise. Then there exists a type t_a so that $(x_1z_1, t_a) \in R_a^1(\mathcal{T})$. 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 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 $(x_1z_1, t_a) \in R_a^2(\mathcal{T})$. Thus, $Q_a^2 \neq \text{proj}_{S_a} R_a^2(\mathcal{T})$. ◊

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 Equation (1): 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.

(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 next see that an analog of Proposition 6.2 does not hold for 3-BRSs.

EXAMPLE 6.2. Return to the game in Figure 2. Consider (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

$$Q_a^3 \times Q_b^3 = Q_a^2 \times \{y_1q_1, y_2\}.$$

probability 1 to y_3; and x_2 is a unique sequential best response under a CPS that assigns probability 1 to $\{y_1q_1, y_1q_2\}$. Second, y_1q_1 (resp. y_1q_2) and y_2 are the *only* strategies that are 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_1z_2 (resp. x_1z_1) conditional on observing x_1. Third, y_2 is a unique sequential best response under a CPS that assigns probability 1 to x_1z_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, $x_1 z_2 \notin \text{proj}_{S_a} R_a^2(\mathcal{T})$. ◊

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_b^3 that is a sequential best response under a CPS that strongly believes Q_a^2. But that CPS cannot strongly believe $\text{proj}_{S_a} R_a^2(\mathcal{T})$; this arises because Q_a^2 is a strict subset of $\text{proj}_{S_a} R_a^2(\mathcal{T})$.

Let us review what led to the situation where $Q_a^2 \subsetneq \text{proj}_{S_a} R_a^2(\mathcal{T})$. 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 arise.

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 nontrivial, 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 $[s_c] \subseteq \mathbb{BR}[\text{p}_c]$.[8]

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. Here, "unique" is taken to mean "up to equivalent strategies." Section 8 discusses which games are generic.

By restricting attention to generic games, we solve the problem that arose in Examples 6.1 and 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 occurs despite the fact that $y_1 q_1$ is 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

[8]In BoSOO, *Out-L* and *Out-R* are two equivalent strategies. We have simply been writing *Out*; our notation formally describes an equivalence class of strategies.

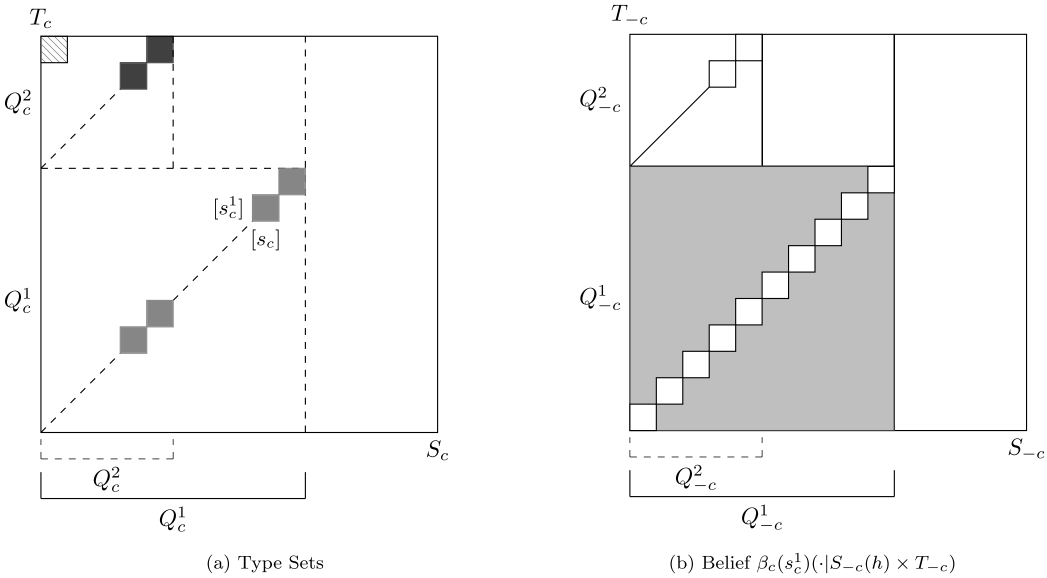

FIGURE 3. Construction of type structure.

p_b^*. (In this game form, $[y_1q_1] = \{y_1q_1\}$.) 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_1z_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 and t_b with $\beta_a(t_a)((y_2, t_b)|S_b \times T_b) = 1$ and $\text{marg}_{S_a} \beta_b(t_b) = p_b^*$.) This would solve the problem seen in Examples 6.1 and 6.2.

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

THEOREM 6.1. *Suppose* Γ *is generic. The following hold for each* m.

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

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

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 $n = 0, \dots, m-1$, 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). (The proof can be found in Appendix B.) Suppose Γ is generic and fix a 2-BRS (Q^0, Q^1, Q^2). The goal is to construct a type structure $\mathcal{T}$ so that $\text{proj}_S R^1(\mathcal{T}) = Q^1$ and $\text{proj}_S R^2(\mathcal{T}) = Q^2$.

Figure 3a 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. Specifically, 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 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 3a, the diagonal of $Q_c^1 \times Q_c^1$ is the union over gray boxes along the pictorial diagonal of $Q_c^1 \times Q_c^1$. The diagonal of $Q_c^2 \times Q_c^2$ is the union over black 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, $(Q_c^1 \times Q_c^1) \setminus \text{diag}_c^i$).

The idea is to construct belief maps so that $R_c^1(\mathcal{T})$ is contained in the union of squares $(Q_c^1 \times Q_c^1) \cup (Q_c^2 \times Q_c^2)$ and $R_c^2(\mathcal{T})$ 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, we will ask that the following properties hold:

(DIAG.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.

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

$$\text{diag}_c^2 \subseteq R_c^2(\mathcal{T}) \subseteq \big(Q_c^2 \times Q_c^2\big) \qquad \text{and} \qquad \text{diag}_c^1 \subseteq R_c^1(\mathcal{T}) \setminus R_c^2(\mathcal{T}) \subseteq \big(Q_c^1 \times Q_c^1\big).$$

We may well have $\text{diag}_c^2 \subsetneq R_c^2(\mathcal{T}) \subseteq (Q_c^2 \times Q_c^2)$. That is, pictorially, $R_c^2(\mathcal{T})$ may well contain both the diagonal black boxes and the (off-diagonal) striped box. However, we require that $\text{diag}_c^1 = R_c^1(\mathcal{T}) \setminus R_c^2(\mathcal{T}) \subseteq (Q_c^1 \times Q_c^1)$. With this, each off-diagonal point in $Q_c^1 \times Q_c^1$ is 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 $\text{marg}_{S_{-c}} \beta_c(s_c^2) = j_c(s_c^2)$. 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.[9] Figure 3b illustrates this case; note that the off-diagonal (illustrated by the gray filling) is nonempty. Moreover, the off-diagonal meets each strategy in Q_{-c}^1. (Formally, for each $s_{-c} \in Q_{-c}^1$, $(\{s_{-c}\} \times T_{-c}) \cap ((Q_{-c}^1 \times Q_{-c}^1) \setminus \text{diag}_{-c}^1) \neq \emptyset$.) We can then choose $\beta_c(s_c^1)$ so that (i) $\text{marg}_{S_{-c}} \beta_c(s_c^1) = j_c(s_c^1)$, (ii) for each h, $\beta_c(s_c^1)(S_{-c} \times Q_{-c}^1|S_{-c}(h) \times T_{-c}) = 1$, and

[9]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 strongly believe Q_{-c}^1.

		Bob C	Bob D
Ann	C	κ, κ	γ, ζ
Ann	D	ζ, γ	δ, δ

FIGURE 4. Prisoner's Dilemma.

(iii) for each h, $\beta_c(s_c^1)(\text{diag}_{-c}^1 | S_{-c}(h) \times T_{-c}) = 0$. So each $\beta_c(s_c^1)$ has beliefs that are concentrated on 1-strategy-type pairs, but off the diagonal.

Observe that under the construction,

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

Since the same holds for the other player $-c$, 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 $R_{-c}^1(\mathcal{T})$, while each 2-type s_c^2 strongly believes $R_{-c}^1(\mathcal{T})$. As such,

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

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

7. ANALYZING GAMES

Write $\text{EFR}^m = \text{EFR}_a^m \times \text{EFR}_b^m$ for the set of strategies that survives m rounds of EFR. We have that $(S, \text{EFR}^1, \ldots, \text{EFR}^m)$ is an m-BRS. So EFR^m is contained in S^m. However, the containment may be strict. In particular, there may be an m-BRS $(Q^0, Q^1, \ldots, Q^m)$ with $\text{EFR}^m \cap Q^m = \emptyset$. This can already be seen from BoSOO (Figure 1). There, for each m, $(Out, R) \in S^m$, but $R \notin \text{EFR}_b^2$ and $Out \notin \text{EFR}_a^3$. This can arise because strong belief is non-monotonic. (See Examples 4.1 and 4.2.)

With Theorem 6.1 in mind, we turn to analyzing games via the m-BRS. We provide two examples: the three-repeated Prisoner's Dilemma and the Centipede game (Rosenthal (1981)). Both have a somewhat subtle relationship to EFR. This, in turn, has implications for the interpretation of experimental data.

Repeated Prisoner's Dilemma Consider the Prisoner's Dilemma in Figure 4, where $\zeta > \kappa > \delta > \gamma$ and $2\kappa > \zeta + \gamma$. For notational simplicity, we focus on the three-repeated game. Corollary D.1 establishes that the game is generic. Thus, Theorem 6.1 applies and so we can focus on m-BRSs.

Write V^t for the set of t-period nodes. Each nonterminal node is associated with a sequence of moves; for instance, $(\phi, (C, D))$ is the two-period node that follows Ann playing C in the first period and Bob playing D in the first period. So $V^1 = \{\phi\}$, $V^2 = \{\phi\} \times \{C, D\}^2$, and $V^3 = \{\phi\} \times \{C, D\}^4$. Since, at each node, all past moves are known, we can write a strategy of player c as a mapping $s_c : V \to \{C, D\}$.

We begin with the EFR benchmark. Write $\mathrm{EFR}_c^0 = S_c$. Then, for each $m = 0, 1, 2$,

$$\mathrm{EFR}_c^{m+1} = \big\{s_c \in \mathrm{EFR}_c^m : \text{ for each } v \in V^{3-m} \text{ with } s_c \in S_c(v), s_c(v) = D\big\}.$$

(See Corollary D.1.) For each $m \geq 3$, $\mathrm{EFR}_c^{m+1} = \mathrm{EFR}_c^m$. So if $s_c \in \mathrm{EFR}_c^1 \backslash \mathrm{EFR}_c^2$, then s_c plays D at each three-period node it allows, but it plays C at some two-period node it allows. Likewise, if $s_c \in \mathrm{EFR}_c^2 \backslash \mathrm{EFR}_c^3$, then s_c plays D at each two-period and three-period node it allows, but it plays C at the initial node ϕ.

We know that, for each m, $\mathrm{EFR}^m \subseteq S^m$. In fact, the two sets are equivalent for $m = 1$. However, this is not the case for $m \geq 2$. In particular, the following hold:

PROPOSITION 7.1.

(i) $S_c^1 = \mathrm{EFR}_c^1$.

(ii) S_c^2 *is the set of strategies* $s_c \in \mathrm{EFR}_c^1$ *so that*

(a) $D \in s_c(\{(\phi, (s_c(\phi), C)), (\phi, (s_c(\phi), D))\})$, *and*

(b) if $2\delta > \beta + \gamma$, *then* $s_c(\phi) = C$ *implies* $s_c(\phi, (C, C)) = D$.

(iii) S_c^3 *is the set of strategies* $s_c \in S_c^2$ *so that* $s_c(\phi) = D$.

(iv) For each $m \geq 4$, S_c^m *is the set of strategies* $s_c \in S_c^3$ *with* $s_c(\phi, (D, D)) = D$.

So, S_c^2 requires that a strategy s_c play D at some second-period node that it allows, but it does not require that s_c play D at each second-period node. This, in turn, implies that a strategy in S_c^3 initially plays D. With this, S_c^4 is the set of strategies $s_c \in \mathrm{EFR}_c^1$ that satisfy (i) $s_c(\phi) = D$, and (ii) $s_c(\phi, (D, D)) = D$. But note that $\mathrm{EFR}_c^4 \subsetneq S_c^4$: It may well be that $s_a(\phi, (D, C))$ (resp. $s_b(\phi, (C, D))$) is C.

An implication of Proposition 7.1 is that, for each m, S^m is outcome equivalent to EFR^m. More precisely, (i) for each $(s_a, s_b) \in S^m$, there exists some $(r_a, r_b) \in \mathrm{EFR}^m$ so that $\zeta(s_a, s_b) = \zeta(r_a, r_b)$, and (ii) for each $(s_a, s_b) \in \mathrm{EFR}^m$, $(s_a, s_b) \in S^m$. However, for each $m \geq 2$, $\mathrm{EFR}^m \subsetneq S^m$.

The distinction between S^m and EFR^m has important implications for the inferences that can be made from experimental data. To understand why, focus on the case where $\beta + \gamma \geq 2\delta$. Consider an experimental data set obtained by the direct response method, where we observe Ann play C in the second period. Conventional wisdom would suggest that this indicates Ann cannot both be rational and "reason" that Bob is rational. Indeed, this is the conclusion that EFR would suggest. However, this behavior is indeed consistent with R1SBR. (In fact, if in the first period, Ann and Bob played (D, C), then this behavior is consistent with RmSBR for all m.) Suppose, instead, that the experimental data set is obtained by the strategy method and we observe Ann play C at each (relevant) second-period node (i.e., we observe $s_a(\phi, (s_a(\phi), C)) = s_a(\phi, (s_a(\phi), D)) = C$). In that case, we can conclude that Ann's behavior is indeed inconsistent with R1SBR.

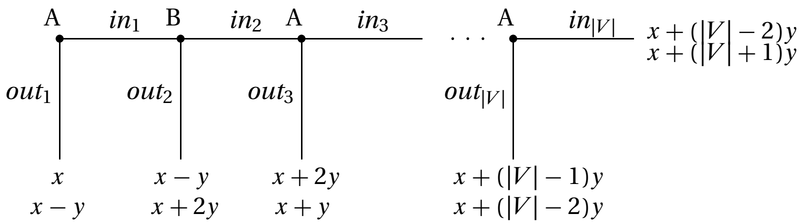

FIGURE 5. Centipede game.

It is, of course, well understood that a benefit of the strategy method is that it provides additional data that may not be observed from the direct response method. (See, e.g., the opening example in Brandts and Charness (2011).) Typically this occurs because Ann has a choice that can only be observed if Bob plays a particular prior action; in the direct response method, Bob may not play the action and so we may not be able to observe Ann's choice. The situation here is different: If we only use the direct response method and we observe Ann choose D in the final period, then we cannot contradict the hypothesis that "Ann is rational and strongly believes Bob is rational" *independent of what Bob plays*. However, if we use the strategy method, we can observe Ann's behavior at each second-period node and so there is potentially observed behavior that would contradict this hypothesis.[10]

Centipede game Figure 5 depicts the Centipede game. We order the nonterminal nodes as $v = 1, 2, \dots, |V|$, where $|V| \geq 3$. (So, $v = 1$ indicates the initial node and $v = |V|$ indicates the last nonterminal 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 $x, y > 0$.[11] 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. $(x + |V|y, x + (|V|-1)y)$). Figure 5 depicts $|V|$ odd.

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 $[out, v]_c$ be the set of strategies of player c that allow v and then play out_v. Likewise, let $[in]_c$ 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 player who moves at node $v \leq |V| - 1$ strictly prefers out_{v+2} (resp. $in_{|V|}$ if $v = |V| - 1$) to out_v and strictly prefers out_v to out_{v+1}. Second, the player who moves at node $|V|$ strictly prefers $out_{|V|}$ to $in_{|V|}$. Third, for each $v \leq |V| - 1$ (resp. $v = |V|$), $[out, v]_c$ is the set of best responses under a CPS that strongly believes $[out, v+1]_{-c}$ (resp. $[in]_{-c}$). Fourth, the game is generic. (This follows from the second and third features.) Thus, Theorem 6.1 applies and, so, we can focus on m-BRSs.[12]

[10] This, of course, presumes that the strategy method works, in that a strategy is optimal in the strategy-method game if and only if it is optimal in the direct-response game. Siniscalchi (2020) provides foundations for the strategy method.

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

[12] Reny's (1993) Take it Or Leave it game also satisfies these properties. Anything we say about Centipede also applies to that game.

A useful benchmark is m rounds of EFR. Observe that

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

Moreover,

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

For all $m \geq |V|+1$, $\mathrm{EFR}^m_\ell \times \mathrm{EFR}^m_{-\ell} = \mathrm{EFR}^{|V|}_\ell \times \mathrm{EFR}^{|V|}_{-\ell}$. Note that 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 7.2. *In the Centipede game, the following hold for each finite $m \geq 1$:*

(i) $S^m_a = \mathrm{EFR}^m_a$.

(ii) If $|V|$ is odd, then $S^m_b = S_b$. If $|V|$ is even, then $S^m_b = (S_b \backslash [in]_b)$.

At first glance, part (i) of Proposition 7.2 may appear trivial: For each m, $\mathrm{EFR}^m_a \times \mathrm{EFR}^m_b$ is consistent with an m-BRS. Thus, $\mathrm{EFR}^m_a \subseteq S^m_a$. However, the key is to show that $S^m_a \subseteq \mathrm{EFR}^m_a$ and, as we have seen, this is not the case for the second mover, Bob. (Appendix D.2 explains why this is the case.)

Proposition 7.2 points to a distinction between the first mover and the second mover—one that is important for interpreting experimental data. Consider the case where $|V|$ is odd, so that the first mover is the last mover. First, for each m, $S^m_a = \mathrm{EFR}^m_a$. So, for instance, if we observe the first mover play $out_{|V|-2}$, then we can conclude that the first mover's behavior is consistent with R3SBR but inconsistent with R4SBR. Second, for each m, $S^m_b = S_b$. So, in particular, any strategy that we observe the second mover play is consistent with RmSBR for each m. (In fact, any strategy that we observe the second mover play is consistent with RCSBR.) This contradicts the conventional wisdom that observing the second player choose *in* indicates that the second player exhibits some form of "bounded reasoning about rationality."

To understand the difference between the first mover and the second mover, note that when $|V|$ is odd, $[out, 1]_a \times S_b$ is an EFBRS. Thus, $S_b \subseteq S^\infty_b$. But, for any nonempty EFBRS $Q_a \times Q_b$, we have $Q_a = [out, 1]_a$.[13] Thus, $S^\infty_a = [out, 1]_a$.

8. Discussion

[13]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 in_1-⋯-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}.

Generic games Theorem 6.1 shows that, in generic games, m-BRSs characterize the R$(m-1)$SBR sequences. This raises the question: "Which games are generic?" We begin with no relevant ties.

DEFINITION 8.1 (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.

We may have a game that satisfies NRT that is nongeneric. (See Example E.1.) However, there is a subclass of NRT games that are generic—ones in which a strategy that is a best response under some CPS is a best response under a "degenerate CPS." (See Definition E.2 and Proposition E.1.) Perfect-information games satisfy that condition and, so, a perfect-information game satisfying NRT is generic. (See Proposition E.2.)

There is a related condition that ensures genericity. 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 (i) $\sigma(r_c) > 0$; and (ii) for all $s_{-c} \in X_{-c} \cap S_{-c}(h)$, $\sum_{r_c \in S_c(h)} \pi_c(r_c, s_{-c})\sigma(r_c) = \pi_c(s_c, s_{-c})$. If $s_c \notin S_c(h)$, then no r_c supports s_c with respect to (X_{-c}, h).

DEFINITION 8.2. A 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. Corollary E.1 (Appendix E.2) establishes that a game that satisfies NRC is generic.

Termination of the m-BRS 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 $Q^{m+1} \subseteq Q^m$, $(Q^0, Q^1, Q^2, \ldots)$ defines an iterative elimination procedure.[14] We refer to this as an m-BRS elimination procedure. Note that 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 analyst knew at which M this occurred, they could use that fact to determine that the elimination procedure has stopped.

At first glance, there may appear to be a 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 $Q^{m+1} \subsetneq Q^m = Q^{m-1}$. To see this, refer to the simultaneous-move game given by Figure 6. For each m, there is an m-BRS

[14] Importantly, an m-BRS is not an order of elimination of iterated conditional dominance (Chen and Micali (2012)).

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

FIGURE 6. Pause is not termination.

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 $(m+1)$. Since m can be any number, we can have arbitrarily long pauses before shrinkage. To understand why this can occur, note that we can have $R^{m+1}(\mathcal{T}) \subsetneq R^m(\mathcal{T})$, but $\text{proj}_S R^{m+1}(\mathcal{T}) = \text{proj}_S R^m(\mathcal{T})$. When this happens, it may well be that $\text{proj}_S R^{m+2}(\mathcal{T}) \subsetneq \text{proj}_S R^{m+1}(\mathcal{T})$.

Despite this, we can provide a bound on the elimination procedure $(S^0, S^1, S^2, \ldots)$, i.e., we can find some M so that, for all $m \geq M$, $S^m = S^M$. 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 $Q_c^{m+1} \backslash Q_c^{m+2}$—must be contained in 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 8.1. *If $(Q^0, \ldots, Q^m)$ is an m-BRS with $Q^m = Q^{m-1}$, then Q^{m-1} is an EFBRS.*

Fix some $(Q^0, Q^1, Q^2, \ldots)$ where (i) for each m, $(Q^0, \ldots, Q^m)$ is an m-BRS, and (ii) $Q^{n+1} = Q^n$. Then Q^n is an EFBRS. We can define a new sequence $(\hat{Q}^0, \hat{Q}^1, \hat{Q}^2, \ldots)$ so that (i) for each $m \leq n$, $\hat{Q}^m = Q^m$, and (ii) for each $m > n$, $\hat{Q}^m = Q^n$. Then, for each m, $(\hat{Q}^0, \ldots, \hat{Q}^m)$ is an m-BRS. So, for each m, $Q^m \subseteq \hat{Q}^m \subseteq S^m$. From this, we get the following termination result.

PROPOSITION 8.1. *Set*

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

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

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

In practice, it is often not necessary to compute all the M-BRSs. Refer to Figure 7. Begin with $Q^0 = S$ and identify all the 1-BRSs (Q^0, Q^1). Use these 1-BRSs to identify all the 2-BRSs (Q^0, Q^1, Q^2). And so on. Notice that along any given M-BRS path $(Q^0, Q^1, \ldots, Q^M)$, we can stop at $m < M$ if $Q^m = Q^{m+1}$.

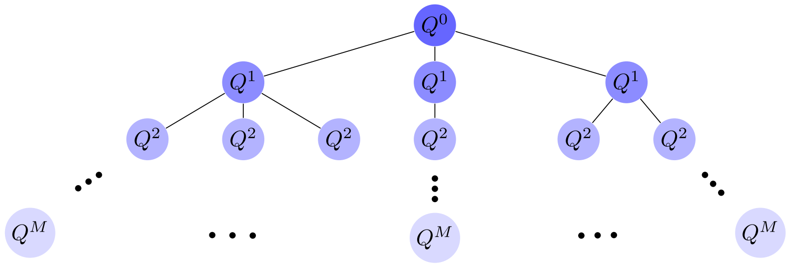

FIGURE 7. The m-BRS elimination tree.

Computing an m-BRS: NRC games Games that satisfy NRC have a simple characterization of the m-BRS concept, one that allows for a "simpler" computation of the m-BRS concept.

PROPOSITION 8.2. *Suppose Γ satisfies NRC. Then $(Q^0, \ldots, Q^m)$ forms an m-BRS if and only if*

(i) Q^1 is nonempty, and

(ii) for each $n = 1, \ldots, m$ and each $s_c \in Q^n_c$, there exists $\mathrm{p}_c \in \mathcal{A}(S_{-c}, \mathcal{E}_c)$ with $s_c \in \mathbb{BR}[\mathrm{p}_c]$ and p_c strongly believes $Q^0_{-c}, \ldots, Q^{n-1}_{-c}$.

Fix 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. If the game satisfies NRC, Proposition 8.2 provides two ways that simplify making that determination. First, we can replace CPSs with arrays. Second, we can eliminate the maximality criterion. Appendix E.5 explains why each of these simplifications may fail absent NRC. It also shows that when we reduce the definition in this way, repeated application of a simplex algorithm determines whether or not $(Q^0, \ldots, Q^m)$ is an m-BRS.

Beyond generic games It would be desirable to have a procedure that determines the sets S^m in nongeneric 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), and (BRP.3), and does not strongly believe Q^1_b. However, under that amendment, we loose an analog of Proposition 6.1; for a given $\mathcal{T}$, $(\mathrm{proj}_S R^0(\mathcal{T}), \ldots, \mathrm{proj}_S R^m(\mathcal{T}))$ 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 $Q^n \subseteq \mathrm{proj}_S R^n(\mathcal{T})$ for each $n = 1, \ldots, m$. The example leaves open the possibility that there may be an alternate m-BRS $(\hat{Q}^0, \hat{Q}^1, \ldots, \hat{Q}^m)$ so that the following hold: (i) $\hat{Q}^m = Q^m$, and (ii) there exists some type structure $\mathcal{T}$ so that $\hat{Q}^n \subseteq \mathrm{proj}_S R^n(\mathcal{T})$ 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.

Finite-order type structures Type structures induce hierarchies of conditional beliefs about the strategies played, i.e., mth-order beliefs for all m. This suggests that players can contemplate sentences of the form "I think that you think that I think ...." As such, one might incorrectly hypothesize that our analysis requires that players have an unlimited ability to engage in interactive reasoning, despite the fact that they exhibit "bounded reasoning about rationality" (formalized as RmSBR, but not R$(m+1)$SBR). However, there is no such requirement. The key observation is that hierarchies of beliefs beyond level m do not affect R$(m-1)$SBR. Formally, consider two types t_a and u_a with the same mth-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. Thus, the higher-order beliefs become a consequence of our formalism and do not have implications for our characterization result. In particular, we could instead adapt the finite-order type structure approach proposed in Kets (2010) and Heifetz and Kets (2018)—amended for CPSs—and 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.)

Heifetz–Kets rationalizability Heifetz and Kets (2018) define a notion of rationalizability for finite-order settings. It is quite different from the analysis here. They focus on simultaneous-move games of incomplete information and use a finite-order type structure to model incomplete information. Strategic uncertainty is captured implicitly by their rationalizability concept, which we call HK rationalizability. Importantly, their notion of rationalizability is different from the standard notion (to which ours reduces in simultaneous-move games). If we apply their concept to a game of complete information, a strategy may survive two rounds of HK rationalizability even though it does not survive two rounds of Bernheim (1984) and Pearce (1984) rationalizability. This arises for the same reason that a level-2 type in the cognitive hierarchy model (Camerer et al. (2004)) may not play a 2-rationalizable strategy. A level-2 type in the cognitive hierarchy model assigns positive probability to a level-0 type and, so, may assign positive probability to a dominated strategy.[15]

Appendix A: Preliminaries

This appendix provides preliminary lemmas that are used in subsequent results.

Marginalization property of belief

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

Proof. Suppose $\beta_c(t_c)$ strongly believes $E_{-c} \subseteq S_{-c} \times T_{-c}$. 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

$$\begin{aligned}\text{marg}_{S_{-c}} \beta_c\big(\text{proj}_{S_{-c}} E_{-c}|S_{-c}(h)\big) &= \beta_c\big(\text{proj}_{S_{-c}} E_{-c} \times T_{-c}|S_{-c}(h) \times T_{-c}\big) \\ &\geq \beta_c\big(E_{-c}|S_{-c}(h) \times T_{-c}\big) = 1.\end{aligned}$$

[15]This does not occur in the level-k model (Costa-Gomes et al. (2001)) and, as such, a level-2 type in the level-k model does survive two rounds of rationalizability.

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 CPSs Fix a CPS $\mathrm{p}_c \in \mathcal{C}(S_{-c}, \mathcal{E}_c)$ 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

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

We refer to q_c as the *image CPS of* p_c *under* τ_{-c}. So defined, q_c is indeed a CPS. See Battigalli et al. (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 $\mathrm{p}_c \in \mathcal{C}(S_{-c}, \mathcal{E}_c)$, we can find some CPS $q_c \in \mathcal{C}(S_{-c} \times T_{-c}, \mathcal{E}_c \otimes T_{-c})$ so that $\text{marg}_{S_{-c}} q_c = p_c$.

Structure of games and sequential best responses By perfect recall, we have the following properties. (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.

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

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

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

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

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

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

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,

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

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, it is not the case that $S(h) \subseteq S(h^{**})$. 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 and 6.2

Proof of Proposition 6.1. The proof is by induction on m.

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

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

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

Proof of Proposition 6.2. Begin with part (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 $\operatorname{marg}_{S_c} \beta_c(s_c)$ 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 $\operatorname{proj}_{S_c} R^1_c(\mathcal{T}) = Q^1_c$.

Turn to part (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 $\mathbb{BR}[j_c(s_c)] \subseteq Q^2_c$.

With this in mind, set $T_c = Q_c^1$ and define $\beta_c(s_c)$ so that $\mathrm{marg}_{S_{-c}} \beta_c(s_c) = j_c(s_c)$. 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

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

Then

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

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

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 an effective singleton, then we simply say it is *nonsingleton*.

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 $j_c^n : Q_c^n \to \mathcal{C}(S_{-c}, \mathcal{E}_c)$ that satisfy the following criteria:

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

(j.2) 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 $j_c^n : Q_c^n \to \mathcal{C}(Q_{-c}, \mathcal{E}_c)$ that satisfy condition (j.2). But, condition (j.1) is stronger than that required by an m-BRS. If we find mappings $j_c = (j_c^1, \ldots, j_c^m)$ that satisfy these requirements, 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* $(Q^0, \ldots, Q^m)$ *that satisfies the strong justification property. Then there exists a type structure* $\mathcal{T}$ *so that, for each* $n = 1, \ldots, m$, $\mathrm{proj}_S R^n(\mathcal{T}) = Q^n$.

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

We now turn to prove these two results.

C.1 *Proof of Proposition C.1*

Throughout we fix an m-BRS $(Q^0, \ldots, Q^m)$ that satisfies 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 is $T_c = \bigsqcup_{n=1}^m U_c^n$. We 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 is

$$\mathrm{diag}_c^n = \bigcup_{s_c \in Q_c^n} \big([s_c] \times v_c^n([s_c])\big).$$

Observe that, if $[s_c] = [r_c]$, then $v_c^n([s_c]) = v_c^n([r_c])$ and so $[s_c] \times v_c^n([r_c]) \subseteq \mathrm{diag}_c^n$. Moreover, if Q_c^n is nonsingleton, then, for each $s_c \in Q_c^n$, there exists a type $t_c \in U_c^n$ so that $(s_c, t_c) \in (Q_c^n \times U_c^n) \setminus \mathrm{diag}_c^n$.

For each $n = 1, \ldots, m$, 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 satisfy the following: For $n = 1$, if Q_{-c}^1 is nonsingleton, 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 $\tau_{-c}^1(s_{-c}) \in (S_{-c} \times U_{-c}^1) \setminus \mathrm{diag}_{-c}^1$. For $n = 2, \ldots, m$ and 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 $\tau_{-c}^n(s_{-c}) = (s_{-c}, v_{-c}^\ell(s_{-c}))$.

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

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

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

By conditions (j.1) and (j.2) of strong justification, $\mathrm{diag}_c^n \subseteq \mathbb{Q}_c^n$.

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

Proof. If $s_c \in Q_c^n$, then $s_c \in \mathbb{BR}[j_c^n(s_c)]$ and so $(s_c, v_c^n(s_c)) \in \mathbb{Q}_c^n$. Fix some $(s_c, v_c^n(r_c)) \in \mathbb{Q}_c^n$. 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.

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

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

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

$n = 2$: Fix some $k = 1, \ldots, m$ and some $(r_c, v_c^k(s_c)) \in \mathbb{BR}[j_c^k(s_c)] \times \{v_c^k(s_c)\} \subseteq \mathbb{Q}_c^k$. We show that (i) and (ii) hold. To do so, we make use of the following properties: $R_{-c}^1(\mathcal{T}) = \bigcup_{k=1}^m \mathbb{Q}_{-c}^k$ and $Q_{-c}^1 = \mathrm{proj}_{S_{-c}} \bigcup_{k=1}^m \mathbb{Q}_{-c}^k = \mathrm{proj}_{S_{-c}} R_{-c}^1(\mathcal{T})$ (Lemma C.1).

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

Second, suppose that $k=1$ and Q^1_{-c} is nonsingleton. Observe that, in this case,

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

By condition (j.1) 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 $(s_{-c}, t_{-c}) \notin R^1_{-c}(\mathcal{T})$. Thus, $v^1_c(s_c)$ does not strongly believe $R^1_{-c}(\mathcal{T})$.

Finally, suppose that $k = 2, \ldots, m$. Fix a conditioning event $S_{-c}(h) \times T_{-c}$ so that $R^1_{-c}(\mathcal{T}) \cap (S_{-c}(h) \times T_{-c}) \neq \emptyset$. Since $Q^1_{-c} = \text{proj}_{S_{-c}} \mathbb{Q}^1_{-c} = \text{proj}_{S_{-c}} R^1_{-c}(\mathcal{T})$, 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,

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

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

$n \geq 3$: 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) and (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 $R^{n-2}_c(\mathcal{T})$. Thus, $v^k_c(s_c)$ does not strongly believe $R^{n-1}_{-c}(\mathcal{T})$.

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

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

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

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

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, $R^{n-1}_{-c}(\mathcal{T}) = \bigcup_{l=n-1}^{m} \mathbb{Q}^l_{-c}$, it follows that $\beta_c(v^k_c(s_c))(R^{n-1}_{-c}(\mathcal{T}) | S_{-c}(h) \times T_{-c}) = 1$, as desired. $\square$

The proof of Proposition C.1 is immediate from Lemmas C.1 and 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 $s_c \in \mathbb{BR}[\mathrm{p}_c]$. 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 $[s^*_c] = \mathbb{BR}[\mathrm{p}_c]$ and p_c does not strongly believe $[s^*_{-c}]$.*

To show the lemma, it is 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,*

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

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 $s_c, r_c \in S_c(h)$,

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

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 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 will establish 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

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

The fact that we can choose r^*_c in this way follows from Lemma A.3.[16]

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 $S_c(h)$.

[16]Specifically: Let $\bar{H}^0_c$ 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 $h^1 \in \bar{H}^0_c$ and note that $r^*_c \in S_c(h^1)$. Choose r^1_c to solve Equation (4) for $h = h^1$ and set $r^*_c(h) = r^1_c(h)$. Then define $\bar{H}^1_c$ to be the set $h \in \bar{H}^0_c$ so that $r^1_c \in S_c(h)$ 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 $\bar{H}^K_c = \emptyset$ has been constructed. Then "fill in" $r^*_c(h)$ arbitrarily at all information sets h for which it has not been defined. (Note that r^*_c precludes those information sets.)

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 3.2 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 3.2, $p_c(s_{-c}|S_{-c}(h)) > 0$ implies $p_c(s_{-c}|S_{-c}(h^*)) > 0$. By construction, for any s_{-c} with $p_c(s_{-c}|S_{-c}(h^*)) > 0$, $\zeta(r_c^*, s_{-c}) = \zeta(r_c, s_{-c})$.

Let $\alpha = p_c(S_{-c}(h) \backslash S_{-c}(h^*)|S_{-c}(h))$. 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

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

Note that

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

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

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

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

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

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 v 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 $v \in h$ 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^ε_c so that

$$q^\varepsilon_c(s^*_{-c}|S_{-c}) = 1 - \varepsilon \quad \text{and} \quad q^\varepsilon_c(r^*_{-c}|S_{-c}) = \varepsilon,$$

and, for each $h \in H_c$ with $S_{-c}(h) \cap \{s^*_{-c}, r^*_{-c}\} = \emptyset$, $q^\varepsilon_c(\cdot|S_{-c}(h)) = p_c(\cdot|S_{-c}(h))$. Note that the unique CPS q^ε_c 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

$$\begin{aligned}&\sum_{s_{-c} \in S_{-c}} [\pi_c(s^*_c, s_{-c}) - \pi_c(r_c, s_{-c})] q^\varepsilon_c(s_{-c}|S_{-c}) \\ &\quad = (1-\varepsilon)[\pi_c(s^*_c, s^*_{-c}) - \pi_c(s_c, s^*_{-c})] + \varepsilon[\pi_c(s^*_c, r^*_{-c}) - \pi_c(s_c, r^*_{-c})] > 0\end{aligned}$$

for each $r_c \in S_c(h)$. Thus, $\mathbb{BR}[\mathrm{q}^\varepsilon_c] = [s^*_c]$ 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^ε_c so that

$$q^\varepsilon_c(s_{-c}|S_{-c}) = \begin{cases} 1 - \varepsilon & \text{if } s_{-c} = s^*_{-c}, \\ \varepsilon \mu^*_c(s_{-c}) & \text{if } s_{-c} \neq s^*_{-c}. \end{cases}$$

and, for each $h \in H_c$, with $S_{-c} \cap (\{s^*_{-c}\} \cup \operatorname{Supp} \mu^*_c) = \emptyset$, $q^\varepsilon_c(\cdot|S_{-c}(h)) = p_c(\cdot|S_{-c}(h))$. Note that the unique CPS q^ε_c that satisfies these conditions does not strongly believe $[s^*_{-c}]$. We show that we can choose $\varepsilon > 0$ so that $\mathbb{BR}[\mathrm{q}^\varepsilon_c] = [s^*_c]$. There are three steps.

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 $\varepsilon \in (0, \bar{\varepsilon}(r_c))$,

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

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 $\varepsilon \in (0, \bar{\varepsilon}(r_c))$,

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

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

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

As such,

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

for all $\varepsilon > 0$.

Step 2. Take $\bar{\varepsilon} = \min\{\bar{\varepsilon}(r_c) : r_c \in S_c\}$. We 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 follows.

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, $q_c^\varepsilon(\operatorname{Supp}\mu_c^*|S_{-c}(h)) = 0$ 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

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

(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,

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

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

Appendix D: Analyzing games

D.1 *Repeated Prisoner's Dilemma*

It is convenient to adopt the following notational convention: Fix some player c and write $(\phi, (\alpha_c^1, \alpha_{-c}^1)) \in V^2$ for a second-period history (resp. $(\phi, (\alpha_c^1, \alpha_{-c}^1), (\alpha_c^2, \alpha_{-c}^2)) \in V^3$ for a third-period history). Note that when we write $(\alpha_c^t, \alpha_{-c}^t)$, we refer to a vector that first specifies a t-period action of player c and then specifies a t-period action of player $-c$. So when we fix some player c and look at the history $(\phi, (C, D))$, we mean the history where, in the first period, player c chooses C and player $-c$ chooses D. Write $\widehat{\mathrm{EFR}}_c^1$ for the set of s_c so that $s_c(v) = D$ for each $v \in V^3$ with $s_c \in S_c(v)$.

D.1.1 *Unique best responses* We begin by recording results about best responses. This serves three purposes. First, we will use the results to establish that the game is generic. Second, we will use the results to establish that $\widehat{\mathrm{EFR}}_c^1 = \mathrm{EFR}_c^1$, a claim made in the text. Third, we will use the results to prove Proposition 7.1.

Lemma D.1. *Let $s_c^* \in \widehat{\mathrm{EFR}}_c^1$ be such that one of the following conditions holds:*

(i) $s_c^*(\phi) = C$ *and* $s_c^*(\phi, (C, C)) = D$*;*

(ii) $s_c^*(\phi) = D$ *and* $s_c^*(\phi, (D, D)) = D$*;*

(iii) $s_c^*(\phi) = D$, $s_c^*(\phi, (D, D)) = C$, *and* $s_c^*(\phi, (D, C)) = D$*; or*

(iv) if $\zeta + \gamma > 2\delta$, $s_c^*(\phi) = C$, $s_c^*(\phi, (C, C)) = C$, *and* $s_c^*(\phi, (C, D)) = D$.

Then there exists some $\mathrm{p}_c \in \mathcal{C}(S_{-c}, \mathcal{E}_c)$ *and some* $s_{-c}^* \in \widehat{\mathrm{EFR}}_{-c}^1$ *so that* $\mathbb{BR}[\mathrm{p}_c] = [s_c^*]$ *and* p_c *strongly believes* $[s_{-c}^*]$.

Lemma D.2. *Let $s_c^* \in \widehat{\mathrm{EFR}}_c^1$ be such that one of the following conditions holds:*

(i) $s_c^*(\phi) = C$ *and* $s_c^*(\phi, (C, C)) = C$*; or*

(ii) $s_c^*(\phi) = D$, $s_c^*(\phi, (D, C)) = C$, *and* $s_c^*(\phi, (D, D)) = C$.

Then there exists some $\mathrm{p}_c \in \mathcal{C}(S_{-c}, \mathcal{E}_c)$ *so that* $\mathbb{BR}[\mathrm{p}_c] = [s_c^*]$.

As a corollary of these lemmas we have the following result.

Corollary D.1.

(i) The three-repeated Prisoner's Dilemma is generic.

(ii) $\mathrm{EFR}_c^1 = \widehat{\mathrm{EFR}}_c^1$.

Proof. It is immediate that $\mathrm{EFR}_c^1 \subseteq \widehat{\mathrm{EFR}}_c^1$. Thus, both parts follow if, for each strategy $s_c \in \widehat{\mathrm{EFR}}_c^1$, there is a CPS p_c so that $\mathbb{BR}[\mathrm{p}_c] = [s_c]$. This follows from Lemma D.1(i)–(iii) and Lemma D.2(i) and (ii). □

The proofs of Lemmas D.1 and D.2 proceed as follows. We begin by fixing a strategy s^*_{-c} and showing that $\pi_c(s^*_c, s^*_{-c}) \geq \pi_c(r_c, s^*_{-c})$ for each $r_c \in \widehat{\mathrm{EFR}}^1_c$ with $\zeta(s^*_c, s^*_{-c}) \neq \zeta(r_c, s^*_{-c})$. Next, we observe that there is exactly one $v^* \in V^2$ so that $s^*_c \in S_c(v^*)$ and $s^*_{-c} \notin S_{-c}(v^*)$. (This follows since each player has two actions.) We then fix a strategy $r^*_{-c} \in S_{-c}(v^*)$ and show that, for each $r_c \in \widehat{\mathrm{EFR}}^1_c \cap S_c(v^*)$ with $\zeta(s^*_c, r^*_{-c}) \neq \zeta(r_c, r^*_{-c})$, $\pi_c(s^*_c, r^*_{-c}) \geq \pi_c(r_c, r^*_{-c})$. Moreover, for each $r_c \in \widehat{\mathrm{EFR}}^1_c \backslash [s^*_c]$, either $\pi_c(s^*_c, s^*_{-c}) > \pi_c(r_c, s^*_{-c})$ or $\pi_c(s^*_c, r^*_{-c}) > \pi_c(r_c, r^*_{-c})$.[17] Now consider a CPS p_c so that

(i) $p_c(\{s^*_{-c}\} \mid S_{-c}(v)) = 1$, if $s^*_{-c} \in S_{-c}(v)$; and

(ii) $p_c(\{r^*_{-c}\} \mid S_{-c}(v)) = 1$, if $s^*_{-c} \notin S_{-c}(v)$ but $r^*_{-c} \in S_{-c}(v)$.

(All other choices are arbitrary.) Then $\mathbb{BR}[\mathrm{p}_c] = [s^*_c]$ and p_c strongly believes $[s^*_{-c}]$. In the specific case of Lemma D.1, we choose $s^*_{-c} \in \widehat{\mathrm{EFR}}^1_{-c}$, allowing for a stronger conclusion.

Remark D.1. There is one case of interest not covered by Lemma D.1: namely $\zeta + \gamma = 2\delta$. In that case, there is one other strategy—called r_c in the proof—so that $\pi(s^*_c, s^*_{-c}) = \pi(r_c, s^*_{-c})$ and $r_c \notin S_c(v^*)$. In that case, under the construction, $\mathbb{BR}[\mathrm{p}_c] = [s^*_c] \cup [r_c]$. We will make use of that fact.

We now complete the non-common features of the proofs.

Proof of Lemma D.1(i). Let s^*_{-c} be a strategy with $s^*_{-c}(\phi) = C$, $s^*_{-c}(\phi, (C, C)) = C$, $s^*_{-c}(\phi, (D, C)) = D$, and, for each $v \in V^3$, $s^*_{-c}(v) = D$. Note that $s^*_{-c} \in \widehat{\mathrm{EFR}}^1_{-c}$. Observe that $\pi_c(s^*_c, s^*_{-c}) = \kappa + \zeta + \delta$. Fix some $r_c \in \widehat{\mathrm{EFR}}^1_c$ with $\zeta(r_c, s^*_{-c}) \neq \zeta(s^*_c, s^*_{-c})$. There are three possible cases:

- $r_c(\phi) = D$ and $r_c(\phi, (D, C)) = C$: $\pi_c(r_c, s^*_{-c}) = \zeta + \gamma + \delta$,
- $r_c(\phi) = D$ and $r_c(\phi, (D, C)) = D$: $\pi_c(r_c, s^*_{-c}) = \zeta + 2\delta$, and
- $r_c(\phi) = C$ and $r_c(\phi, (C, C)) = C$: $\pi_c(r_c, s^*_{-c}) = 2\kappa + \delta$.

In each case, $\pi_c(s^*_c, s^*_{-c}) > \pi_c(r_c, s^*_{-c})$.

There is a single history $v^* \in V^2$ so that $s^*_c \in S_c(v^*)$ but $s^*_{-c} \notin S_{-c}(v^*)$, namely $v^* = (\phi, (C, D))$. We will choose distinct strategies r^*_{-c} based on whether $s^*_c(v^*) = C$ or $s^*_c(v^*) = D$.

First, suppose $s_c(v^*) = C$. Let r^*_{-c} be such that $r^*_{-c}(\phi) = D$ and, for each $v \in V^2 \cup V^3$, $r^*_{-c}(v) = C$ if and only if $v = (\cdot, (\alpha^{t-1}_c, \alpha^{t-1}_{-c}))$ with $\alpha^{t-1}_c = C$. Then $\pi_c(s^*_c, r^*_{-c}) = \gamma + \kappa + \zeta$. Fix $r_c \in \widehat{\mathrm{EFR}}^1_c \cap S_c(v^*)$ with $\zeta(r_c, r^*_{-c}) \neq \zeta(s^*_c, r^*_{-c})$. Thus, $r_c(v^*) = D$ and so $\pi_c(r_c, r^*_{-c}) = \gamma + \zeta + \delta < \pi_c(s^*_c, r^*_{-c})$.

Second, suppose $s_c(v^*) = D$. Let r^*_{-c} be such that $r^*_{-c}(v) = D$ for each $v \in V$. Then $\pi_c(s^*_c, r^*_{-c}) = \gamma + 2\delta$. Fix $r_c \in \widehat{\mathrm{EFR}}^1_c \cap S_c(v^*)$ with $\zeta(r_c, r^*_{-c}) \neq \zeta(s^*_c, r^*_{-c})$. Thus, $r_c(v^*) = C$ and so $\pi_c(r_c, r^*_{-c}) = 2\gamma + \delta < \pi_c(s^*_c, r^*_{-c})$. □

[17] A detail in how this is implemented in the proof: For each $r_c \in \widehat{\mathrm{EFR}}^1_c \backslash [s^*_c]$, either $\zeta(s^*_c, s^*_{-c}) \neq \zeta(r_c, s^*_{-c})$ or $\zeta(s^*_c, r^*_{-c}) \neq \zeta(r_c, r^*_{-c})$. This again follows from the fact that each player has two actions.

Remark D.2. The proof of Lemma D.1(i) establishes that we can take s^*_{-c} so that $s^*_{-c}(\phi)=C$, $s^*_{-c}(\phi,(C,C))=C$, $s^*_{-c}(\phi,(D,C))=D$, and, for each $v\in V^3$, $s^*_{-c}(v)=D$.

Proof of Lemma D.1(ii). Let s^*_{-c} be a strategy with $s^*_{-c}(v)=D$ for each $v\in V$. Note that $s^*_{-c}\in\widehat{\mathrm{EFR}}^1_{-c}$. Observe that $\pi_c(s^*_c,s^*_{-c})=3\delta$. Fix some $r_c\in\widehat{\mathrm{EFR}}^1_c$ with $\zeta(r_c,s^*_{-c})\neq\zeta(s^*_c,s^*_{-c})$. Then either (i) $r_c(\phi)=C$, or (ii) $r_c(\phi)=D$, and $r_c(\phi,(D,D))=C$. As such, $\pi_c(r_c,s^*_{-c})\leq 2\delta+\gamma<\pi_c(s^*_c,s^*_{-c})$.

There is a single history $v^*\in V^2$ so that $s^*_c\in S_c(v^*)$ but $s^*_{-c}\notin S_{-c}(v^*)$, namely $v^*=(\phi,(D,C))$. We will choose distinct strategies r^*_{-c} based on whether $s^*_c(v^*)=C$ or $s^*_c(v^*)=D$.

First, suppose $s_c(v^*)=C$. Let r^*_{-c} be such that (i) $r^*_{-c}(\phi)=C$, (ii) for each $v\in V^2$, $r^*_{-c}(v)=C$ if and only if $v=(\phi,(D,\cdot))$, and (iii) for each $v\in V^3$, $r^*_{-c}(v)=C$ if and only if $v=(\cdot,(\cdot,\cdot),(C,\cdot))$. Note that $\pi_c(s^*_c,r^*_{-c})=2\zeta+\kappa$. Fix $r_c\in\widehat{\mathrm{EFR}}^1_c\cap S_c(v^*)$ with $\zeta(r_c,r^*_{-c})\neq\zeta(s^*_c,r^*_{-c})$. Then $r_c(v^*)=D$ and so $\pi_c(r_c,r^*_{-c})=2\zeta+\delta<\pi_c(s^*_c,r^*_{-c})$.

Second, suppose $s_c(v^*)=D$. Let r^*_{-c} be such that (i) $r^*_{-c}(\phi)=C$, and (ii) for each $v\in V^2\cup V^3$, $r^*_{-c}(v)=D$. Note that $\pi_c(s^*_c,r^*_{-c})=\zeta+2\delta$. Fix $r_c\in\widehat{\mathrm{EFR}}^1_c\cap S_c(v^*)$ with $\zeta(r_c,r^*_{-c})\neq\zeta(s^*_c,r^*_{-c})$. Then $r_c(v^*)=C$ and so $\pi_c(r_c,r^*_{-c})=\zeta+\gamma+\delta<\pi_c(s^*_c,r^*_{-c})$. □

Remark D.3. The proof of Lemma D.1(ii) establishes that we can take s^*_{-c} so that $s^*_{-c}(v)=D$ for each $v\in V$.

Proof of Lemma D.1(iii). Let s^*_{-c} be a strategy with $s^*_{-c}(\phi)=C$ and, for each $v\in V^2\cup V^3$, $s^*_{-c}(v)=D$. Note that $s^*_{-c}\in\widehat{\mathrm{EFR}}^1_{-c}$. Observe that $\pi_c(s^*_c,s^*_{-c})=\zeta+2\delta$. Fix some $r_c\in\widehat{\mathrm{EFR}}^1_c$ with $\zeta(r_c,s^*_{-c})\neq\zeta(s^*_c,s^*_{-c})$. There are three possible cases:

- $r_c(\phi)=D$ and $r_c(\phi,(D,C))=C$: $\pi_c(r_c,s^*_{-c})=\zeta+\gamma+\delta$,
- $r_c(\phi)=C$ and $r_c(\phi,(C,C))=D$: $\pi_c(r_c,s^*_{-c})=\gamma+2\delta$, and
- $r_c(\phi)=C$ and $r_c(\phi,(C,C))=C$: $\pi_c(r_c,s^*_{-c})=2\gamma+\delta$.

In each case, $\pi_c(s^*_c,s^*_{-c})>\pi_c(r_c,s^*_{-c})$.

There is a single history $v^*\in V^2$ so that $s^*_c\in S_c(v^*)$ but $s^*_{-c}\notin S_{-c}(v^*)$, namely $v^*=(\phi,(D,D))$. Let r^*_{-c} be such that (i) $r^*_{-c}(\phi)=D$, (ii) for each $v\in V^2$, $r^*_{-c}(v)=C$ if and only if $v=(\phi,(D,\cdot))$, and (iii) for each $v\in V^3$, $r^*_{-c}(v)=C$ if and only if $v=(\cdot,(\cdot,\cdot),(C,\cdot))$. Note that $\pi_c(s^*_c,r^*_{-c})=\delta+\kappa+\zeta$. Fix $r_c\in\widehat{\mathrm{EFR}}^1_c\cap S_c(v^*)$ with $\zeta(r_c,r^*_{-c})\neq\zeta(s^*_c,r^*_{-c})$. Then $r_c(v^*)=D$ and so $\pi_c(r_c,r^*_{-c})=2\delta+\zeta<\pi_c(s^*_c,r^*_{-c})$. □

Remark D.4. The proof of Lemma D.1(iii) establishes that we can take s^*_{-c} so that $s^*_{-c}(v)=D$ for each $v\in V^2\cup V^3$.

Proof of Lemma D.1(iv). Let s^*_{-c} be a strategy with (i) $s^*_{-c}(\phi)=D$, (ii) for each $v\in V^2$, $s^*_{-c}(v)=C$ if and only if $v=(\phi,(C,\cdot))$, and (iii) for each $v\in V^3$, $s^*_{-c}(v)=D$. Note that $s^*_{-c}\in\widehat{\mathrm{EFR}}^1_{-c}$. Observe that $\pi_c(s^*_c,s^*_{-c})=\gamma+\zeta+\delta$. Fix some $r_c\in\widehat{\mathrm{EFR}}^1_c$ with $\zeta(r_c,s^*_{-c})\neq\zeta(s^*_c,s^*_{-c})$. There are three possible cases:

- $r_c(\phi) = D$ and $r_c(\phi, (D, C)) = C$: $\pi_c(r_c, s^*_{-c}) = 2\delta + \gamma$,
- $r_c(\phi) = D$ and $r_c(\phi, (D, C)) = D$: $\pi_c(r_c, s^*_{-c}) = 3\delta$, and
- $r_c(\phi) = C$ and $r_c(\phi, (C, D)) = C$: $\pi_c(r_c, s^*_{-c}) = \gamma + \kappa + \delta$.

In each case, $\pi_c(s^*_c, s^*_{-c}) > \pi_c(r_c, s^*_{-c})$. (Here we use the fact that $\gamma + \zeta > 2\delta$.)

There is a single history $v^* \in V^2$ so that $s^*_c \in S_c(v^*)$ but $s^*_{-c} \notin S_{-c}(v^*)$, namely $v^* = (\phi, (C, C))$. Let r^*_{-c} be such that (i) $r^*_{-c}(\phi) = C$, (ii) for each $v \in V^2$, $r^*_{-c}(v) = C$ if and only if $v = (\phi, (C, \cdot))$, and (iii) for each $v \in V^3$, $r^*_{-c}(v) = C$ if and only if $v = (\cdot, (\cdot, \cdot), (C, \cdot))$. Note that $\pi_c(s^*_c, r^*_{-c}) = 2\kappa + \zeta$. Fix $r_c \in \widehat{\mathrm{EFR}}^1_c \cap S_c(v^*)$ with $\zeta(r_c, r^*_{-c}) \neq \zeta(s^*_c, r^*_{-c})$. Then $r_c(v^*) = D$ and so $\pi_c(r_c, r^*_{-c}) = \kappa + \zeta + \delta < \pi_c(s^*_c, r^*_{-c})$. □

Remark D.5. Fix $s^*_c(\phi) = C$, $s^*_c(\phi, (C, C)) = C$, and $s^*_c(\phi, (C, D)) = D$. Suppose $\gamma + \zeta = 2\delta$. Then the proof of Lemma D.1(iv) establishes that there exists some CPS p_c and some $s^*_{-c} \in \widehat{\mathrm{EFR}}^1_{-c}$ so that $[s^*_c] \subseteq \mathbb{BR}[\mathrm{p}_c]$ and p_c strongly believes $[s^*_{-c}]$.

Proof of Lemma D.2 (i). Construct a strategy s^*_{-c} with $s^*_{-c}(\phi) = C$ and, for each $v \in V^2 \cup V^3$, $s^*_{-c}(v) = C$ if and only if $v = (\cdot, (\alpha^{t-1}_c, \alpha^{t-1}_{-c}))$ with $\alpha^{t-1}_c = C$. Observe that $\pi_c(s^*_c, s^*_{-c}) = 2\kappa + \zeta$. Fix some $r_c \in \widehat{\mathrm{EFR}}^1_c$ with $\zeta(r_c, s^*_{-c}) \neq \zeta(s^*_c, s^*_{-c})$. There are three possible cases:

- $r_c(\phi) = D$ and $r_c(\phi, (D, C)) = C$: $\pi_c(r_c, s^*_{-c}) = 2\zeta + \gamma$,
- $r_c(\phi) = D$ and $r_c(\phi, (D, C)) = D$: $\pi_c(r_c, s^*_{-c}) = \zeta + 2\delta$, and
- $r_c(\phi) = C$ and $r_c(\phi, (C, C)) = D$: $\pi_c(r_c, s^*_{-c}) = \kappa + \zeta + \delta$.

In each case, $\pi_c(s^*_c, s^*_{-c}) > \pi_c(r_c, s^*_{-c})$. (Here we use the assumption that $2\kappa > \zeta + \gamma$.)

There is a single history $v^* \in V^2$ so that $s^*_c \in S_c(v^*)$ but $s^*_{-c} \notin S_{-c}(v^*)$, namely $v^* = (\phi, (C, D))$. We will choose distinct strategies r^*_{-c} based on whether $s^*_c(v^*) = C$ or $s^*_c(v^*) = D$.

First, suppose $s_c(v^*) = C$. Let r^*_{-c} be a strategy so that $r^*_{-c}(\phi) = D$ and, for each $v \in V^2 \cup V^3$, $r^*_{-c}(v) = s^*_{-c}(v)$. Then $\pi_c(s^*_c, r^*_{-c}) = \gamma + \kappa + \zeta$. Fix $r_c \in \widehat{\mathrm{EFR}}^1_c \cap S_c(v^*)$ with $\zeta(r_c, r^*_{-c}) \neq \zeta(s^*_c, r^*_{-c})$. Thus, $r_c(v^*) = D$ and so $\pi_c(r_c, r^*_{-c}) = \gamma + \zeta + \delta < \pi_c(s^*_c, r^*_{-c})$.

Second, suppose $s_c(v^*) = D$. Let r^*_{-c} be a strategy so that, for each $v \in V$, $r^*_{-c}(v) = D$. Then $\pi_c(s^*_c, r^*_{-c}) = \gamma + 2\delta$. Fix $r_c \in \widehat{\mathrm{EFR}}^1_c \cap S_c(v^*)$ with $\zeta(r_c, r^*_{-c}) \neq \zeta(s^*_c, r^*_{-c})$. Thus, $r_c(v^*) = C$ and so $\pi_c(r_c, r^*_{-c}) = 2\gamma + \delta < \pi_c(s^*_c, r^*_{-c})$. □

Proof of Lemma D.2(ii). Construct a strategy s^*_{-c} with (i) $s^*_{-c}(\phi) = C$, (ii) $s^*_{-c}(\phi, (\alpha^1_c, \cdot)) = C$ if and only if $\alpha^1_c = D$, and (iii) $s^*_{-c}(\cdot, (\alpha^2_c, \cdot)) = C$ if and only if $\alpha^2_c = C$. Observe that $\pi_c(s^*_c, s^*_{-c}) = 2\zeta + \kappa$. Fix some $r_c \in \widehat{\mathrm{EFR}}^1_c$ with $\zeta(r_c, s^*_{-c}) \neq \zeta(s^*_c, s^*_{-c})$. There are three possible cases:

- $r_c(\phi) = D$ and $r_c(\phi, (D, C)) = D$: $\pi_c(r_c, s^*_{-c}) = 2\zeta + \delta$,
- $r_c(\phi) = C$ and $r_c(\phi, (C, C)) = C$: $\pi_c(r_c, s^*_{-c}) = \kappa + \gamma + \zeta$, and

- $r_c(\phi) = C$ and $r_c(\phi, (C, C)) = D$: $\pi_c(r_c, s^*_{-c}) = 2\kappa + \zeta$.

In each case, $\pi_c(s^*_c, s^*_{-c}) > \pi_c(r_c, s^*_{-c})$.

There is a single history $v^* \in V^2$ so that $s^*_c \in S_c(v^*)$ but $s^*_{-c} \notin S_{-c}(v^*)$, namely $v^* = (\phi, (D, D))$. Let r^*_{-c} be a strategy so that (i) $r^*_{-c}(\phi) = D$, (ii) $r^*_{-c}(\phi, (\alpha^1_c, \cdot)) = C$ if and only if $\alpha^1_c = D$, and (iii) $r^*_{-c}(\cdot, (\alpha^2_c, \cdot)) = C$ if and only if $\alpha^2_c = C$. Then $\pi_c(s^*_c, r^*_{-c}) = \delta + \kappa + \zeta$. Fix $r_c \in \widehat{\mathrm{EFR}}^1_c \cap S_c(v^*)$ with $\zeta(r_c, r^*_{-c}) \neq \zeta(s^*_c, r^*_{-c})$. Then $r_c(v^*) = D$ and so $\pi_c(r_c, r^*_{-c}) = 2\delta + \zeta < \pi_c(s^*_c, r^*_{-c})$. □

D.1.2 *m-BRSs* Corollary D.1 establishes part (i) of Proposition 7.1: $S^1_c = \mathrm{EFR}^1_c = \widehat{\mathrm{EFR}}^1_c$. To show part (ii), it will be useful to have the following lemmas.

Lemma D.3. *Fix a strategy $s_c \in \mathrm{EFR}^1_c$ so that*

(i) $D \in s_c(\{(\phi, (s_c(\phi), C)), (\phi, (s_c(\phi), D))\})$; and

(ii) if $2\delta > \zeta + \gamma$, then $s_c(\phi) = C$ implies $s_c(\phi, (C, C)) = D$.

Then there exists some 2-BRS (Q^0, Q^1, Q^2) so that $s_c \in Q^2_c \subseteq S^2_c$.

Proof. Fix a strategy s_c as in the statement of the lemma. By Corollary D.1, $s_c \in \widehat{\mathrm{EFR}}^1_c$. Moreover, by Lemma D.1, Remark D.5, and Corollary D.1, there exists some $s_{-c} \in \widehat{\mathrm{EFR}}^1_{-c} = \mathrm{EFR}^1_{-c}$ and a CPS p_c that strongly believes $[s_{-c}]$ so that $[s_c] \subseteq \mathbb{BR}[\mathrm{p}_c]$. Construct $Q^1_c = Q^2_c = \mathbb{BR}[\mathrm{p}_c]$, $Q^1_{-c} = [s_{-c}]$, and $Q^2_{-c} = \emptyset$. Applying Lemmas D.1 and D.2, (S, Q^1) is a 1-BRS. As such, (S, Q^1, Q^2) is a 2-BRS with $s_c \in Q^2_c \subseteq S^2_c$. □

Lemma D.4. *Fix an $(m+1)$-BRS $(Q^0, \dots, Q^m, Q^{m+1})$, where $m \geq 1$. If $v \in V^2$ with $Q^m_{-c} \cap S_{-c}(v) \neq \emptyset$, then, for each $s_c \in Q^{m+1}_c \cap S_c(v)$, $s_c(v) = D$.*

Proof. Fix $v \in V^2$ with $Q^m_{-c} \cap S_{-c}(v) \neq \emptyset$ and $s_c \in Q^{m+1}_c \cap S_c(v)$. Then there exists $\mu \in \Delta(S_c(v))$ with (i) $\pi(s_c, \mu) \geq \pi(r_c, \mu)$ for each $r_c \in S_c(v)$, and (ii) $\mu(Q^m_{-c}) = 1$. Note that, since $m \geq 1$, each $s_{-c} \in Q^m_{-c} \cap S_{-c}(v)$ has $s_{-c}(v') = D$ for all $v' \in V^3$ with $s_{-c} \in S_{-c}(v')$. Thus, $s_c(v) = D$. □

Proof of Proposition 7.1(ii). Lemma D.3 implies that each strategy specified in the result is contained in S^2_c. So we focus on the converse.

Fix a strategy $s_c \in \mathrm{EFR}^1_c = \widehat{\mathrm{EFR}}^1_c$. Suppose there is a 2-BRS, viz. (Q^0, Q^1, Q^2), with $s_c \in Q^2_c$. This implies that there exists some CPS p_c so that $s_c \in \mathbb{BR}[\mathrm{p}_c]$ and p_c strongly believes Q^1_{-c}. Thus, $Q^1_{-c} \neq \emptyset$. As such, there is some $v \in V^2$ with $s_c \in S_c(v)$ and $Q^1_{-c} \cap S_{-c}(v) \neq \emptyset$. By Lemma D.4, $s_c(v) = D$. So either $s_c(\phi, (s_c(\phi), C)) = D$ or $s_c(\phi, (s_c(\phi), D)) = D$.

For the remainder of the proof, suppose that $s_c(\phi) = C$, $s_c(\phi, (C, C)) = C$, and $s_c(\phi, (C, D)) = D$. It suffices to show that, for each $s_{-c} \in Q^1_{-c}$, $s_{-c}(\phi) = D$. If so, then

$$\sum_{s_{-c} \in S_{-c}} \pi_c(s_c, s_{-c}) p\big(\{s_{-c}\} \mid S_{-c}\big) \leq \gamma + \zeta + \delta.$$

But there is an alternate strategy r_c with

$$\sum_{s_{-c}\in S_{-c}} \pi_c(r_c, s_{-c})p\big(\{s_{-c}\} \mid S_{-c}\big) \geq 3\delta.$$

(In particular, take $r_c(v) = D$ for all $v \in V$.) Since $s_c \in \mathbb{BR}[\mathrm{p}_c]$, $\zeta + \gamma \geq 2\delta$.

We now show that, for each $s_{-c} \in Q^1_{-c}$, $s_{-c}(\phi) = D$: Since s_c is a best response under $p_c(\cdot \mid S_{-c}(\phi, (C, C)))$, there must be some s_{-c} and $v \in V^3$ with $p_c(\{s_{-c}\} \mid S_{-c}(\phi, (C, C))) > 0$, $s_{-c}(v) = C$, and $s_{-c} \in S_{-c}(v)$. This implies that $p_c(Q^1_{-c} \mid S_{-c}(\phi, (C, C))) < 1$. Given that p_c strongly believes Q^1_{-c}, $Q^1_{-c} \cap S_{-c}(\phi, (C, C)) = \emptyset$; that is, each $s_{-c} \in Q^1_{-c}$ specifies $s_{-c}(\phi) = D$. □

Lemma D.5. *Fix a strategy* $s_c \in S^2_c$ *with* $s_c(\phi) = D$ *and* $s_c(\phi, (D, D)) = D$. *Then, for each* m, $s_c \in S^m_c$.

Proof. Let s_{-c} be such that $s_{-c}(v) = D$ for all $v \in V$. Lemma D.1 and Remark D.3 imply that $[s_c] \times [s_{-c}]$ is an EFBRS. Define $Q^1 = \cdots = Q^m = [s_c] \times [s_{-c}]$ and note that $(S, Q^1, \ldots, Q^m)$ is an m-BRS. □

Proof of Proposition 7.1(iii). Fix a 3-BRS (Q^0, Q^1, Q^2, Q^3) and some $s_c \in Q^3_c$. Then there exists some CPS p_c so that $[s_c] \subseteq \mathbb{BR}[\mathrm{p}_c]$ and p_c strongly believes Q^2_{-c}. Suppose, contra hypothesis, that $s_c(\phi) = C$. Then $Q^2_{-c} \neq \emptyset$. Moreover, for each $s_{-c} \in Q^2_{-c}$ and each $v \in V^2 \cup V^3$ with $(s_c, s_{-c}) \in S(v)$, $s_c(v) = s_{-c}(v) = D$. (See Lemma D.4.) So, for each $s_{-c} \in Q^2_{-c}$, either $\pi_c(s_c, s_{-c}) = \kappa + 2\delta$ or $\pi_c(s_c, s_{-c}) = \gamma + 2\delta$. Since $p(Q^2_{-c} \mid S_{-c}) = 1$, there exists some $q \in [0, 1]$ with

$$\sum_{s_{-c}\in S_{-c}} \pi_c(s_c, s_{-c})p\big(\{s_{-c}\} \mid S_{-c}\big) = q\kappa + (1-q)\gamma + 2\delta.$$

Consider an alternate strategy r_c with $r_c(v) = D$ for each $v \in V$. Observe that

$$\sum_{s_{-c}\in S_{-c}} \pi_c(r_c, s_{-c})p\big(\{s_{-c}\} \mid S_{-c}\big) \geq q\zeta + (1-q)\delta + 2\delta.$$

This contradicts $[s_c] \subseteq \mathbb{BR}[\mathrm{p}_c]$.

Now fix some $s_c \in S^2_c$ with $s_c(\phi) = D$. We show that there exists a 3-BRS (Q^0, Q^1, Q^2, Q^3) with $s_c \in Q^3_c$. If $s_c(\phi, (D, D)) = D$, this follows from Lemma D.5. So suppose that $s_c(\phi, (D, D)) = C$. It is convenient to define strategies r_c and s_{-c}. In particular, take $r_c \in \widehat{\mathrm{EFR}}^1_c$ so that (i) $r_c(\phi) = C$, (ii) $r_c(\phi, (C, C)) = C$, and (iii) $r_c(\phi, (C, D)) = D$. Take $s_{-c} \in \widehat{\mathrm{EFR}}^1_{-c}$ so that (a) $s_{-c}(\phi) = C$ and (b) for each $v \in V^2 \cup V^3$, $s_{-c}(v) = D$. Set $Q^1_c = [s_c] \cup [r_c]$ and $Q^2_c = Q^3_c = [s_c]$. Set $Q^1_{-c} = Q^2_{-c} = [s_{-c}]$ and $Q^3_{-c} = \emptyset$. By Lemmas D.1 and D.2 and Remark D.2, (S, Q^1, Q^2, Q^3) is a 3-BRS. □

The proof of Proposition 7.1(iv) is immediate from Lemmas D.4 and D.5.

D.2 *Centipede*

Throughout this subsection, fix an m-BRS $(Q^0, Q^1, \dots, Q^m)$ of the Centipede game. We will show that $Q_a^m \subseteq \mathrm{EFR}_a^m$. Begin with the following observation.

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

Lemma D.6. *One of the following must hold:*

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

(ii) $[out, |V|]_\ell \cap Q_\ell^1 = \emptyset$ *and* $|V| = 3$.

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

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

$$\{s_\ell : s_\ell(v) = out_v \text{ for some } v \le |V| - 2\}.$$

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

Lemma D.7. *Fix some* $m = 3, \dots, |V| - 1$. *If* m *is odd, then either*

(i) $[out, |V| + 3 - m]_\ell \cap Q_\ell^m = \emptyset$, *or*

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

If m *is even, then either*

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

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

Proof. We show the base cases of $m = 3, 4$. The inductive step simply repeats those arguments up to relabeling. Note that, since $|V| - 1 \ge m \ge 3$, $|V| \ge 4$. So, by Lemma D.6, $[in]_{-\ell} \cap Q_{-\ell}^2 = \emptyset$. We repeatedly use this fact below.

$m = 3$. Throughout, we suppose that $[out, |V|]_\ell \subseteq Q_\ell^1$. (If not, then we are done.) From this, Lemma D.6 gives that $[in]_{-\ell} \cap Q_{-\ell}^2 = \emptyset$. We divide the argument into two cases.

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

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

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

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

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

(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 $p_\ell([out, \tilde{v}+1]_{-\ell}|S_{-\ell}(\tilde{v})) = 1$. Thus, at node $\tilde{v}$, $[out, \tilde{v}]_\ell$ is a unique best response. So certainly $[out, |V|]_\ell \cap Q^3_\ell = \emptyset$.

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

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

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

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

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

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

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

Corollary D.2. *If $|V| = m$, then either $Q^{|V|}_a = [out, 1]_a$ or $Q^{|V|}_a = \emptyset$.*

Proof. We show the result for $|V|$ odd. (The case of $|V|$ even is analogous.) If $Q^{|V|-2}_a \notin \{[out, 1]_a, \emptyset\}$, the claim is immediate. So, suppose otherwise. By Observation D.1, $[in]_a \notin Q^{|V|-2}_a$. By Lemma D.6, for each $m \leq |V|-2$ odd, $[out, |V|+3-m]_a \cap Q^{|V|-2}_a = \emptyset$. So,

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

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

Appendix E: Proofs for Section 8

E.1 *Canonical CPS*

We begin with a mathematical step useful in several results below. 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.

Fix a strategy s_c^* and some array $\mathrm{p}_c \in \mathcal{A}(S_{-c}, \mathcal{E}_c)$ with $s_c^* \in \mathbb{BR}[\mathrm{p}_c]$. 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 that $h^0 \in \overline{H}_c^0$.

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.

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

Property 2: There is no $h \in H_c^{k+1}$ so that $S_{-c}(h^{k+1}) \subsetneq S_{-c}(h)$.

Property 3: 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 $h^0, h^1, \ldots, h^K$. Refer to these as *basic information sets*. (Note that they depend on both p_c and s_c^*.) We set $q_c(\cdot|S_{-c}(h^k))$ to coincide with the original array $p_c(\cdot|h^k)$. 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.

Lemma E.1. *Fix a strategy s_c^* and some array $\mathrm{p}_c \in \mathcal{A}(S_{-c}, \mathcal{E}_c)$ with $s_c^* \in \mathbb{BR}[\mathrm{p}_c]$. Let $\mathrm{q}_c \in \mathcal{C}(S_{-c}, \mathcal{E}_c)$ be the canonical CPS for (s_c^*, p_c). 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}.

To prove Lemma E.1, it will be useful to have the following lemma.

Lemma E.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 2 of the construction, $S_{-c}(h^k) = S_{-c}(h)$. So $S_c(h^k) \subsetneq S_c(h)$. But then, by Property 3 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 $h \in \overline{H}_c^k$, then $\emptyset \neq S_{-c}(h) \subseteq S_{-c}(h^k)$. It follows from Lemma A.2 that $S_c(h) \cap S_c(h^k) = \emptyset$, a contradiction. □

Proof of Lemma E.1. First, we show that $s_c^* \in \mathbb{BR}[\mathrm{q}_c]$. (That implies $[s_c^*] \subseteq \mathbb{BR}[\mathrm{q}_c]$.) Toward that end, fix some $h \in H_c$ with $s_c^* \in S_c(h)$. Observe that there exists a k such that $h \in \overline{H}_c^k$, i.e., there exists a basic h^k such that $q_c(\cdot|S_{-c}(h))$ is derived from $p_c(\cdot|S_{-c}(h^k))$ by conditioning. (Note that 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 Lemmas E.2 and A.3 that s_c^* is optimal under $q_c(\cdot|S_{-c}(h))$ given all strategies in $S_c(h)$.

Second, we show that if p_c strongly believes E_{-c}, then q_c strongly believes E_{-c}. 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 $q(E_{-c}|S_{-c}(h)) = 1$. □

E.2 *Generic games: No relevant ties*

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

Example E.1. The game in Figure 8 satisfies no relevant ties. Yet it is not generic: D is a sequential best response under p_a if and only if $p_a(L|S_b) = p_a(R|S_b) = 1/2$. Thus, $\mathbb{BR}[\mathrm{p}_a] = \{U, M, D\}$ and there is no q_a with $\mathbb{BR}[\mathrm{q}_a] = [D]$. ◊

Note that in Example E.1, D is justifiable, but not optimal under any CPS that involves point beliefs.

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

		B	
		L	R
	U	4, 0	0, 1
A	M	0, 0	4, 1
	D	2,1	2, 0

FIGURE 8. No relevant ties.

So, a CPS is degenerate if each conditional belief is a point belief.

DEFINITION E.2. 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 E.1 is not degenerately justifiable.

PROPOSITION E.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, $r_c \notin \mathbb{BR}[\mathrm{p}_c]$. □

LEMMA E.3. *A perfect-information game satisfying NRT is degenerately justifiable.*

PROOF. Let s_c be a justifiable strategy. Then, by Lemma 1.2.1 in Ben Porath (1997), for each $S_{-c}(h) \in \mathcal{E}_c$ with $s_c \in S_c(h)$, we can find some $s^h_{-c} \in S_{-c}(h)$ so that $\pi_c(s_c, s^h_{-c}) \geq \pi_c(r_c, s^h_{-c})$ for all $r_c \in S_c(h)$. Use the collection $(s^h_{-c} : h \in H_c \cup \{\phi\})$ to form a degenerate array p_c with $s_c \in \mathbb{BR}[\mathrm{p}_c]$. Then, the canonical CPS q_c is degenerate and, by Lemma E.1, $s_c \in \mathbb{BR}[\mathrm{q}_c]$. □

The following result is now immediate from Proposition E.1 and Lemma E.3.

PROPOSITION E.2. *A perfect-information game satisfying no relevant ties is generic.*

E.3 *Generic games: No relevant convexities*

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 that preserves strong belief and, moreover, s^*_c remains a sequential best response. (See Lemma E.1.) We now 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 proves Proposition 8.2 and establishes that a game that satisfies NRC is generic.

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 E.4. *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 E.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 E.5. *Suppose Γ 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 $\hat{p}_c(\cdot|S_{-c}(h)$ among all strategies in $S_c(h)$ if and only if r_c supports s_c^* given $(\operatorname{Supp} p_c(\cdot|S_{-c}(h)), h)$, and (b) $\operatorname{Supp} \hat{p}_c(\cdot|S_{-c}(h)) = \operatorname{Supp} p_c(\cdot|S_{-c}(h))$. (See Lemmas D.2–D.4 in Brandenburger et al. (2008).) Requirement (i) follows from the construction and NRC; requirements (ii) and (iii) follow immediately from the construction. □

The proof of Proposition 8.2 is immediate from Lemmas E.4 and E.5.

Corollary E.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 Lemmas E.4 and E.5, 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 8.2 is that, when NRC is satisfied, we can forgo using CPSs 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

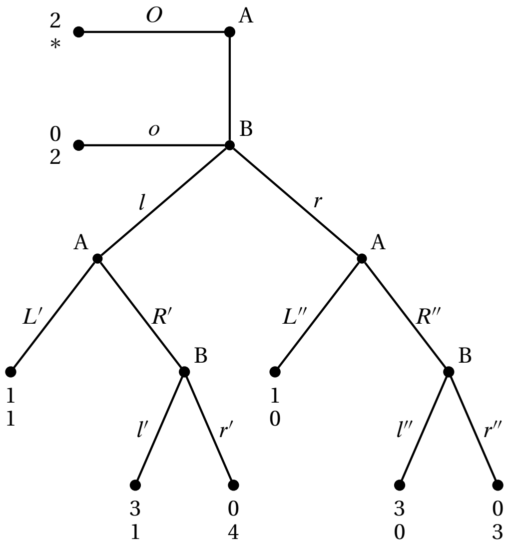

FIGURE 9. A perfect-information game satisfying NRT.

that, for some array $\mathrm{p}_c \in \mathcal{A}(S_{-c}, \mathcal{E}_c)$, conditions (BRP.1), (BRP.2), and (BRP.3) are satisfied. The canonical CPS $\mathrm{q}_c \in \mathcal{C}(S_{-c}, \mathcal{E}_c)$ satisfies conditions (BRP.1) and (BRP.2), but condition (BRP.3) may fail. (An example is available upon request.)

A second implication of Proposition 8.2 is that we can use a weaker maximality criterion. This need not hold in a perfect-information game satisfying NRT, despite the fact that such games are generic.

EXAMPLE E.2. The game in Figure 9 is a perfect-information game satisfying NRT.[18] As such, it is generic. But, the conclusion of Proposition 8.2 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 argue that (Q^0, Q^1, Q^2) satisfies the requirements of Proposition 8.2, but is not a 2-BRS.

Note that O is a unique sequential best response under a CPS that ex ante assigns probability 1 to lr' and 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 on Bob's first information set being reached, the array assigns probability $2/3 : 1/3$ to $L'R'' : R'L''$. So (Q^0, Q^1, Q^2) satisfies the requirements of

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

Proposition 8.2. But, it is not a 2-BRS, since $\mathbb{BR}[\mathrm{p}_b] = \{o, lr', rr''\}$ is not contained in Q^2_b. ◊

E.4 *Termination of the procedure*

Call $\mathcal{Q} = (Q^0, Q^1, \dots)$ a *BR-sequence* if, for each m, $(Q^0, \dots, Q^m)$ is an m-BRS.

Proof of Proposition 8.1. For each BR-sequence $\mathcal{Q} = (Q^0, Q^1, \dots)$, there is a finite $M(\mathcal{Q})$ so that $Q^{M(\mathcal{Q})} = Q^{M(\mathcal{Q})+1}$. Choose the lowest such $M(\mathcal{Q})$ and note that

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

Take 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. $|S_a| = |S_b|$).

Now note that $S^M = S^\infty$: Certainly $S^\infty \subseteq S^M$. Fix $s \in S^M$ and note that there exists some $\mathcal{Q} = (Q^0, Q^1, \dots)$ with $s \in Q^M$ and $Q^M \subseteq Q$ for some EFBRS Q. So $s \in Q \subseteq S^\infty$. □

E.5 *Computing m-BRSs*

Suppose Γ satisfies NRC. Proposition 8.2 offers an alternate approach to computing m-BRSs, one that makes use of arrays. This allows us to use the simplex algorithm to search for appropriate beliefs.

Fix some $(Q^0, \dots, Q^{m-1}, Q^m)$ where $(Q^0, \dots, Q^{m-1})$ is an $(m-1)$-BRS. Also fix some $h \in H_c$ and write $n(h) = \max\{n : Q^n_{-c} \cap S_{-c}(h) \neq \emptyset\}$. Then enumerate

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

Say a strategy $s_c \in Q^m_c$ *passes the test at* h if either $s_c \notin S_c(h)$ or there exist nonnegative numbers $\mu^1, \dots, \mu^K$ with $\sum_{k=1}^K \mu^k = 1$, so that s_c maximizes $\sum_{k=1}^K \pi_c(\cdot, s^k_{-c})\mu^k$ among all strategies in $S_c(h) = \{s^1_c, \dots, s^L_c\}$. 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 whether or not s_c passes the test at h. Specifically, when $s_c \in S_c(h)$, the problem is equivalent to choosing $(\mu^1, \dots, \mu^K, \tau^1, \dots, \tau^L)$ to solve

$$\begin{aligned} \text{maximize} \quad & \sum_{k=1}^K \pi_c(s_c, s^k_{-c})\mu^k \\ \text{subject to} \quad & \sum_{k=1}^K [\pi_c(s_c, s^k_{-c}) - \pi_c(s^l_c, s^k_{-c})]\mu^k + \tau^l = 0 \quad \text{for each } l = 1, \dots, L \\ & \mu^1 + \mu^2 + \cdots + \mu^K = 1 \\ & (\mu^1, \dots, \mu^K, \tau^1, \dots, \tau^L) \ge (0, \dots, 0). \end{aligned}$$

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 et al. (1977).) In the first scenario s_c fails the test; in the latter two scenarios, s_c passes the test.

References

Battigalli, Pierpaolo (1997), "On rationalizability in extensive games." *Journal of Economic Theory*, 74, 40–61. [1626, 1651]

Battigalli, Pierpaolo and Amanda Friedenberg (2009), "Context-dependent forward induction reasoning." http://u.arizona.edu/~afriedenberg/firr-online.pdf. [1611, 1612]

Battigalli, Pierpaolo and Amanda Friedenberg (2012), "Forward induction reasoning revisited." *Theoretical Economics*, 7, 57–98. [1605, 1607, 1608, 1613, 1615]

Battigalli, Pierpaolo, Amanda Friedenberg, and Marciano Siniscalchi (2012), "Strategic uncertainty: An epistemic approach to game theory." (Working Title). [1630]

Battigalli, Pierpaolo and Marciano Siniscalchi (1999), "Hierarchies of conditional beliefs and interactive epistemology in dynamic games." *Journal of Economic Theory*, 88, 188–230. [1611]

Battigalli, Pierpaolo and Marciano Siniscalchi (2002), "Strong belief and forward induction reasoning." *Journal of Economic Theory*, 106, 356–391. [1605, 1607, 1608, 1612, 1613, 1614]

Bernheim, Douglas (1984), "Rationalizable strategic behavior." *Econometrica*, 52, 1007–1028. [1629]

Bradley, Stephen, Arnoldo Hax, and Thomas Magnanti (1977), *Applied Mathematical Programming*. Addison Wesley. [1653]

Brandenburger, Adam (2003), "On the existence of a 'complete' possibility structure." In *Cognitive Processes and Economic Behavior* (Marcello Basili, Nicola Dimitri, and Itzhak Gilboa, eds.), 30–34, Routledge. [1607, 1611]

Brandenburger, Adam, Amanda Friedenberg, and H. Jerome Keisler (2008), "Admissibility in games." *Econometrica*, 76, 307. [1608, 1650]

Brandts, Jordi and Gary Charness (2011), "The strategy versus the direct-response method: A first survey of experimental comparisons." *Experimental Economics*, 14, 375–398. [1624]

Camerer, Colin, Teck-Hua Ho, and Juin-Kuan Chong (2004), "A cognitive hierarchy model of games." *The Quarterly Journal of Economics*, 119, 861–898. [1629]

Chen, Jing and Silvio Micali (2012), "The order independence of iterated dominance in extensive games." *Theoretical Economics*, 8, 125–163. [1626]

Costa-Gomes, Miguel, Vince Crawford, and Bruno Broseta (2001), "Cognition and behavior in normal-form games: An experimental study." *Econometrica*, 69, 1193–1235. [1629]

Friedenberg, Amanda (2010), "When do type structures contain all hierarchies of beliefs?" *Games and Economic Behavior*, 68, 108–129. [1608]

Heifetz, Aviad and Willemien Kets (2018), "Robust multiplicity with a grain of naiveté." *Theoretical Economics*, 13, 415–465. [1629]

Kets, Willemien (2010), "Bounded reasoning and higher-order uncertainty." https://wkets.org. [1629]

Osborne, Martin J. and Ariel Rubinstein (1994), *A Course in Game Theory*. The MIT press. [1610]

Pearce, David (1984), "Rationalizable strategic behavior and the problem of perfection." *Econometrica*, 52, 1029–1050. [1606, 1608, 1613, 1629]

Ben Porath, Elchanan (1997), "Rationality, Nash equilibrium and backwards induction in perfect-information games." *The Review of Economic Studies*, 64, 23–46. [1649]

Reny, Philip (1993), "Common belief and the theory of games with perfect information." *Journal of Economic Theory*, 59, 257–274. [1624]

Rosenthal, Robert (1981), "Games of perfect information, predatory pricing and the chain-store paradox." *Journal of Economic theory*, 25, 92–100. [1609, 1622]

Siniscalchi, Marciano (2020), "Structural rationality in dynamic games." http://faculty.wcas.northwestern.edu/~msi661/. [1624]

Co-editor George J. Mailath handled this manuscript.

Manuscript received 13 May, 2017; final version accepted 27 October, 2020; available online 30 November, 2020.