Theoretical Economics 7 (2012), 57–98 1555-7561/20120057

Forward induction reasoning revisited

Pierpaolo Battigalli
Department of Economics, Bocconi University

Amanda Friedenberg
Department of Economics, Arizona State University

Battigalli and Siniscalchi (2002) formalize the idea of forward induction reasoning as "rationality and common strong belief of rationality" (RCSBR). Here we study the behavioral implications of RCSBR across all type structures. Formally, we show that RCSBR is characterized by a solution concept we call extensive form best response sets (EFBRS's). It turns out that the EFBRS concept is equivalent to a concept already proposed in the literature, namely directed rationalizability (Battigalli and Siniscalchi 2003). We conclude by applying the EFBRS concept to games of interest.

Keywords. Epistemic game theory, forward induction, extensive form best response set, directed rationalizability.

JEL classification. C72.

1. Introduction

Forward induction is a basic concept in game theory. It reflects the idea that players rationalize their opponents' behavior whenever possible. In particular, players form an assessment about the future play of the game, given the information about the past play and the presumption that their opponents are strategic. This affects the players' choices.

Formalizing forward induction reasoning requires an epistemic apparatus: To express the idea that players rationalize their opponents' past behavior, we need a language that explicitly describes what a player believes about the strategies her opponents play and the beliefs they hold at each information set. An (extensive-form based) epistemic type structure gives such a language.

Pierpaolo Battigalli: pierpaolo.battigalli@unibocconi.it
Amanda Friedenberg: amanda.friedenberg@asu.edu
We are indebted to Adam Brandenburger, John Nachbar, and Marciano Siniscalchi for many helpful conversations. Jeff Ely and three referees provided important input—much thanks. We also thank Ethan Bueno de Mesquita, Alfredo Di Tillio, Alejandro Manelli, Elena Manzoni, Andres Perea, Larry Samuelson, Adam Szeidl, and seminar participants at Bocconi University, Boston University, Maastricht University, New York University, Northwestern University, Toulouse, UC Berkeley, the 2009 Southwest Economic Theory Conference, the 2009 North American Econometric Society Meetings, the Kansas Economic Theory Conference, the 2009 SAET conference, and the 2009 European Econometric Society Meetings for important input. Battigalli thanks MIUR and Bocconi University. Friedenberg thanks the W.P. Carey School of Business and the Olin Business School.

DOI: 10.3982/TE598

Within this framework, Battigalli and Siniscalchi (2002) formalize forward induction reasoning using the idea of "strong belief." (See also Stalnaker 1998.) A player *strongly believes* an event E if he assigns probability 1 to E, as long as E is consistent with the information set he has reached. With this, the conditions that each player is rational, strongly believes that "each (other) player is rational," strongly believes "each (other) player is rational and strongly believes others are rational," etc. formally capture the idea of forward induction reasoning. The collection of these assumptions is called *rationality and common strong belief of rationality* (*RCSBR*).

Battigalli and Siniscalchi (2002) analyze the implications of RCSBR in the canonical construction of the so-called universal type structure. (This is a type structure that induces all hierarchies of conditional beliefs.) They show that, in this case, a strategy is consistent with RCSBR if and only if it is extensive-form rationalizable (Pearce 1984). But, for a "smaller" type structure—one that does not induce all hierarchies of conditional beliefs—the strategies consistent with RCSBR may be distinct from the extensive-form rationalizable strategies. (See Battigalli and Siniscalchi 2002 or Example 3 below.)

Given this fact, a natural question arises. What are the implications of forward induction reasoning across all epistemic type structures? The answer is a solution concept we call extensive-form best response sets (EFBRS's). Specifically, we show that RCSBR is characterized by EFBRS's: For a given game and type structure, the strategies consistent with RCSBR form an EFBRS. Conversely, for a given EFBRS, there is a type structure so that the strategies consistent with RCSBR are exactly the given EFBRS. (See Theorem 1.) Of course, the extensive-form rationalizable strategy set is one EFBRS. Which EFBRS obtains depends on the given type structure.

While the EFBRS definition is new, we note that it is equivalent to a definition already proposed in the literature, namely, the directed rationalizability concept. This solution concept is due to Battigalli and Siniscalchi (2003), who refer to it as Δ-rationalizability. We discuss the connection in Section 9.a below. We see that, in some ways, the questions raised here can be viewed as a follow-up to the questions raised in Battigalli and Siniscalchi (2003).

The paper proceeds as follows. The game and epistemic structure are defined in Sections 2 and 3. Rationality and strong belief are defined in Section 4. Section 5 gives the main theorem, a characterization of RCSBR in terms of EFBRS's. Section 6 gives an alternate characterization theorem, in terms of directed rationalizability. We then turn to applications in Sections 7 and 8. Finally, in Section 9, we conclude by discussing certain conceptual and technical aspects of the paper.

2. The game

We consider finite extensive-form games of perfect recall. We write Γ for such a game. The definition we consider is similar to that in Osborne and Rubinstein (1994, Definition 200.1). In particular, it allows for simultaneous moves.[1]

[1] This definition incorporates repeated games. Our analysis does not depend on the specific definition used.

There are two players, namely a (Ann) and b (Bob).[2] Let C_a and C_b be *choice* or *action sets* for Ann and Bob. A history for the game consists of (possibly empty) sequences of simultaneous choices for Ann and Bob. More formally, a *history* is either (i) the empty sequence, written ϕ, or (ii) a sequence of choice pairs $(c^1, \ldots, c^K)$, where $c^k = (c_a^k, c_b^k) \in C_a \times C_b$. Histories have the property that if $(c^1, \ldots, c^K)$ is a history, then so is $(c^1, \ldots, c^L)$ for each $L \le K$. Each history can be viewed as a node in the tree, and so we interchangeably use the terms "node" and "history."

Write x for a history of the game and let $C(x) = \{c \in C_a \times C_b : (x, c)$ is a history for the game$\}$. Write $C_a(x) = \text{proj}_{C_a} C(x)$ and $C_b(x) = \text{proj}_{C_b} C(x)$. By assumption, these sets have the property that $C(x) = C_a(x) \times C_b(x)$. The interpretation is that $C_a(x)$ is the set of *choices available to* a *at history* x. If $|C_a(x)| \ge 2$, say a *moves at history* x or a *is active at* x. (If $|C_a(x)| \le 1$, a is inactive at history x.) Call x a *terminal history* of the game if $C(x) = \varnothing$. (Terminal histories can be viewed either as *terminal nodes* or *paths* for the game.)

Let H_a (resp. H_b) be a partition of the set of all nodes at which a (resp. b) is active plus the initial node ϕ. The partition H_a (resp. H_b) has the property that if x, x' are contained in the same partition member, viz. h in H_a (resp. H_b), then $C_a(x) = C_a(x')$ (resp. $C_b(x) = C_b(x')$). The interpretation is that H_a (resp. H_b) is the family of *information sets* for a (resp. b). (Notice that $\{\phi\} \in H_a \cap H_b$. Perfect recall imposes further requirements on H_a and H_b. See Osborne and Rubinstein 1994, Definition 203.3.) Write $H = H_a \cup H_b$.

Let Z be the set of terminal histories of the game and let z be an arbitrary element of Z. *Extensive-form payoff functions* are given by $\Pi_a : Z \to \mathbb{R}$ and $\Pi_b : Z \to \mathbb{R}$.

We abuse notation and write $C_a(h)$ for the set of choices available to a at information set $h \in H_a$. With this, the set of *strategies* for player a is given by $S_a = \prod_{h \in H_a} C_a(h)$. Define S_b analogously. Each pair of strategies (s_a, s_b) induces a path through the tree. Let $\zeta : S_a \times S_b \to Z$ map each strategy profile into the induced path. *Strategic-form payoff functions* are given by $\pi_a = \Pi_a \circ \zeta$ and $\pi_b = \Pi_b \circ \zeta$. Given a profile (s_a, s_b), write $\pi(s_a, s_b) = (\pi_a(s_a, s_b), \pi_b(s_a, s_b))$ and refer to this payoff vector as an *outcome* of the game. Two strategy profiles, (s_a, s_b) and (r_a, r_b), are *outcome equivalent* if $\pi(s_a, s_b) = \pi(r_a, r_b)$. (Of course, if (s_a, s_b) and (r_a, r_b) induce the same path (i.e., if $\zeta(s_a, s_b) = \zeta(r_a, r_b)$), they are outcome equivalent. But, they may be outcome equivalent even if they do not.)

For each information set $h \in H$, write $S_a(h)$ (resp. $S_b(h)$) for the set of strategies for a (resp. b) that allow h. (That is, $s_a \in S_a(h)$ if there is some $s_b \in S_b$ so that the path induced by (s_a, s_b) passes through h.) Let $\mathcal{S}_a$ (resp. $\mathcal{S}_b$) be the collection of all $S_a(h)$ (resp. $S_b(h)$) for $h \in H_b$ (resp. $h \in H_a$). Thus, $\mathcal{S}_a$ represents the information structure of b about the strategy of a. In particular, at each of b's information sets, he has a belief about a that assigns probability 1 to the set of a's strategies consistent with the information set being reached.

3. The type structure

This section defines an epistemic type structure. There are two ingredients: First, for each player, there are type sets T_a and T_b. Informally, each player "knows" his own type,

[2]The analysis extends to n-player games, up to issues of correlation. See Section 9.b.

but faces uncertainty about the strategy the other player will choose and the type of the other player. So each type $t_a \in T_a$ is associated with a belief on $S_b \times T_b$. Of course, we want to specify a belief at each information set. Therefore, we map each type into a conditional probability system (CPS) on $S_b \times T_b$, where the conditioning events correspond to the information sets in the game tree. That is, for each type, there is an array of probability measures on $S_b \times T_b$, one for each information set, and this array satisfies the rules of conditional probability when possible.

We now give the formal definitions. These closely follow the definitions in Battigalli and Siniscalchi (2002). Throughout, let Ω be a separable metrizable space and let $\mathcal{B}(\Omega)$ be the Borel σ-algebra on Ω. We endow the product of separable metrizable spaces with the product topology and endow a subset of a separable metrizable space with the relative topology. Write $\mathcal{P}(\Omega)$ for the set of Borel probability measures on Ω and endow $\mathcal{P}(\Omega)$ with the topology of weak convergence.

Definition 1 (Rényi 1955). Fix a separable metrizable space Ω and a nonempty collection of events $\mathcal{E} \subseteq \mathcal{B}(\Omega)$. A *conditional probability system* (CPS) on $(\Omega, \mathcal{E})$ is a mapping $\mu(\cdot|\cdot) : \mathcal{B}(\Omega) \times \mathcal{E} \to [0,1]$ such that, for every $E \in \mathcal{B}(\Omega)$ and $F, G \in \mathcal{E}$, the following statements hold:

(i) $\mu(F|F) = 1$,

(ii) $\mu(\cdot|F) \in \mathcal{P}(\Omega)$, and

(iii) $E \subseteq F \subseteq G$ implies $\mu(E|G) = \mu(E|F)\mu(F|G)$.

Call $\mathcal{E}$, with $\varnothing \neq \mathcal{E} \subseteq \mathcal{B}(\Omega)$, a *collection of conditioning events* for Ω.

When it is clear that $\mu(\cdot|\cdot)$ is a CPS on $(\Omega, \mathcal{E})$, we omit reference to its arguments, simply writing μ instead of $\mu(\cdot|\cdot)$.

Write $\mathcal{C}(\Omega, \mathcal{E})$ for the set of conditional probability systems on $(\Omega, \mathcal{E})$. The set $\mathcal{C}(\Omega, \mathcal{E})$ can be viewed as a subset of $[\mathcal{P}(\Omega)]^{\mathcal{E}}$. We endow $[\mathcal{P}(\Omega)]^{\mathcal{E}}$ with the product topology and then endow $\mathcal{C}(\Omega, \mathcal{E})$ with the relative topology. If $\mathcal{E}$ is countable, $\mathcal{C}(\Omega, \mathcal{E})$ is separable metrizable. When the set of conditioning events is clear from the context, we omit reference to $\mathcal{E}$, simply writing $\mathcal{C}(\Omega)$.

We are often interested in product sets. We adopt the convention that if $\Omega_1 \times \Omega_2 = \varnothing$, then both $\Omega_1 = \varnothing$ and $\Omega_2 = \varnothing$. Fix some $\mathcal{E} \subseteq \mathcal{B}(\Omega_1)$ and write $\mathcal{E} \otimes \Omega_2$ for the set of all $E \times \Omega_2$, where $E \in \mathcal{E}$. Of course, $\mathcal{E} \otimes \Omega_2 \subseteq \mathcal{B}(\Omega_1 \times \Omega_2)$.

Consider a CPS $\mu(\cdot|\cdot)$ on $(\Omega_1 \times \Omega_2, \mathcal{E} \otimes \Omega_2)$, where $\mathcal{E} \subseteq \mathcal{B}(\Omega_1)$. Define $\nu(\cdot|\cdot) : \mathcal{B}(\Omega_1) \times \mathcal{E} \to [0,1]$ so that $\nu(E|F) = \mu(E \times \Omega_2 | F \times \Omega_2)$ for all $E \in \mathcal{B}(\Omega_1)$ and $F \in \mathcal{E}$. Then ν is a conditional probability system on $(\Omega_1, \mathcal{E})$. When $\nu(\cdot|\cdot)$ is defined in this way, write $\nu(\cdot|\cdot) = \text{marg}_{\Omega_1} \mu(\cdot|\cdot)$. No confusion should result.

Definition 2. Fix an extensive-form game Γ. A Γ-*based type structure* is a collection

$$\langle S_a, S_b; \mathcal{S}_a, \mathcal{S}_b; T_a, T_b; \beta_a, \beta_b \rangle,$$

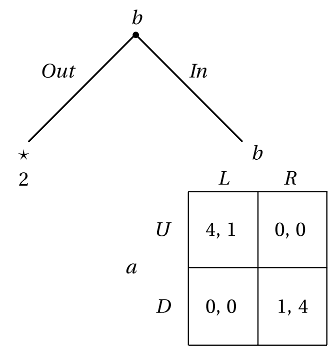

Figure 1. Battle of the sexes with an outside option.

where T_a (resp. T_b) is a nonempty separable metrizable space and $\beta_a : T_a \to \mathcal{C}(S_b \times T_b, \mathcal{S}_b \otimes \mathcal{T}_b)$ (resp. $\beta_b : T_b \to \mathcal{C}(S_a \times T_a, \mathcal{S}_a \otimes \mathcal{T}_a)$) is a measurable *belief map*. Members of T_a (resp. T_b) are called *types*. Members of $S_a \times T_a \times S_b \times T_b$ are called *states*.

To illustrate Definition 2, consider two examples of Γ-based type structures. Each is based on the game Γ of the battle of the sexes (BoS) with an outside option as given in Figure 1.

Example 1. Suppose the game of BoS with an outside option is played in a society that has come to form a "lady's choice convention." Loosely, everyone in the society thinks that if the lady gets to move in a BoS-like situation, she makes choices that can lead to her "best payoff," i.e., she plays *Up*, hoping to get a payoff of 4. Moreover, it is "transparent" that everyone thinks this.

The convention restricts the beliefs players do vs. do not consider possible.[3] It can be modelled by a type structure $\langle S_a, S_b; \mathcal{S}_a, \mathcal{S}_b; T_a, T_b; \beta_a, \beta_b \rangle$ based on the game in Figure 1. The type structure satisfies the following conditions: Each type t_b of Bob is mapped to a CPS on $S_a \times T_a$ that assigns probability 1 to $\{Up\} \times T_a$ at each information set. Moreover, for each such CPS, there is a type of Bob, viz. t_b, so that $\beta_b(t_b)$ is exactly that CPS. Likewise, for each CPS on $S_b \times T_b$, there is a type of Ann, viz. t_a, so that $\beta_a(t_a)$ is exactly that CPS. (See Battigalli and Friedenberg 2009 on how to construct such a structure.)

Notice that at each information set, each type of Bob assigns probability 1 to the event "Ann plays *Up*," i.e., to Ann trying to achieve her best payoff. There are no restrictions on Ann's beliefs about Bob's play of the game. This follows from β_a being onto—for each belief she can have about S_b, there is a type of Ann that has that belief. But at each information set, each type of Ann assigns probability 1 to the event "at each information set, Bob assigns probability 1 to the event 'Ann plays *Up*,'" and so on. In this sense, it is transparent that Bob thinks that if Ann gets to move, she will play *Up*. ◊

Example 2. Suppose the game of BoS with an outside option is played among players who have no reason to believe that the other players are more or less likely to choose

[3]There is no restriction on which strategies the players can vs. cannot play.

a particular strategy or to have particular beliefs, etc. This idea can be modelled by a type structure that contains all possible conditional beliefs (about types), i.e., by a type structure $\langle S_a, S_b; \mathcal{S}_a, \mathcal{S}_b; T_a, T_b; \beta_a, \beta_b \rangle$ based on the game in Figure 1, where β_a and β_b are onto.

This is known as a *complete* type structure. (The terminology is due to Brandenburger 2003.) One example of a complete type structure is the canonical construction of a type structure, as in Battigalli and Siniscalchi (1999a). That type structure induces all hierarchies of conditional beliefs. ◊

4. Rationality and strong belief

We now turn to the main epistemic definitions, all of which have counterparts with a and b reversed. Begin by extending $\pi_a(\cdot,\cdot)$ to $S_a \times \mathcal{P}(S_b)$ in the usual way, i.e., $\pi_a(s_a, \varpi_a) = \sum_{s_b \in S_b} \pi_a(s_a, s_b)\varpi_a(s_b)$. Since the measure ϖ_a on S_b reflects a belief by a about b, we write $\varpi_a \in \mathcal{P}(S_b)$.

Definition 3. Fix $X_a \subseteq S_a$ and $s_a \in X_a$. Say s_a is *optimal under* $\varpi_a \in \mathcal{P}(S_b)$ *given* X_a if $\pi_a(s_a, \varpi_a) \geq \pi_a(r_a, \varpi_a)$ for all $r_a \in X_a$.

Definition 4. Say $s_a \in S_a$ is *sequentially optimal under* $\mu_a(\cdot|\cdot) : \mathcal{B}(S_b) \times \mathcal{S}_b \to [0, 1]$ if, for all h with $s_a \in S_a(h)$, s_a is optimal under $\mu_a(\cdot|S_b(h))$ given $S_a(h)$. Say $s_a \in S_a$ is *sequentially justifiable* if there exists $\mu_a(\cdot|\cdot) : \mathcal{B}(S_b) \times \mathcal{S}_b \to [0, 1]$ so that s_a is sequentially optimal under $\mu_a(\cdot|\cdot)$.

Definition 5. Say (s_a, t_a) is *rational* if s_a is sequentially optimal under $\text{marg}_{S_b} \beta_a(t_a)$.

Let R_a be the set of strategy-type pairs, viz. (s_a, t_a), at which a is rational.

Definition 6 (Battigalli and Siniscalchi 2002). Fix a CPS $\mu(\cdot|\cdot) : \mathcal{B}(\Omega) \times \mathcal{E} \to [0, 1]$ and an event $E \in \mathcal{B}(\Omega)$. Say μ *strongly believes* E if

(i) there exists $F \in \mathcal{E}$ so that $E \cap F \neq \varnothing$ and

(ii) for each $F \in \mathcal{E}$, $E \cap F \neq \varnothing$ implies $\mu(E|F) = 1$.

If a CPS μ strongly believes E and $\Omega \in \mathcal{E}$, then $\mu(E|\Omega) = 1$. In our application, we have $\Omega \in \mathcal{E}$. Of course, no CPS strongly believes the empty set.

Strong belief fails a monotonicity property, i.e., μ may strongly believe an event E but not some event F with $E \subseteq F$. (This can happen if there is some $G \in \mathcal{E}$ with $E \cap G = \varnothing$ but $F \cap G \neq \varnothing$.) But there are two important properties that strong belief does satisfy. (These properties are useful in our analysis.)

Property 1 (Conjunction). *Fix a CPS on $(\Omega, \mathcal{E})$, viz. μ, and a finite or countable collection of events $E_1, E_2, \ldots$. If μ strongly believes $E_1, E_2, \ldots$, then μ strongly believes $\bigcap_m E_m$.*

Property 2 (Marginalization). *Fix a CPS μ on $(\Omega_1 \times \Omega_2, \mathcal{E} \otimes \Omega_2)$, where $\mathcal{E} \subseteq \mathcal{B}(\Omega_1)$. If μ strongly believes $E \in \mathcal{B}(\Omega_1 \times \Omega_2)$ and* $\text{proj}_{\Omega_1} E$ *is Borel, then* $\text{marg}_{\Omega_1} \mu$ *strongly believes* $\text{proj}_{\Omega_1} E$.

Definition 7. Say $t_a \in T_a$ *strongly believes* $E_b \in \mathcal{B}(S_b \times T_b)$ if $\beta_a(t_a)$ strongly believes E_b.

Let $\text{SB}_a(E_b)$ be the set of strategy-type pairs (s_a, t_a) such that t_a strongly believes event E_b. That is, $\text{SB}_a(E_b)$ is the event that "Ann strongly believes E_b."

Now, we inductively define the set of states at which there is rationality and mth-order strong belief of rationality. Set $R_a^1 = R_a$ (resp. $R_b^1 = R_b$). The event that Ann is rational and Ann strongly believes "Bob is rational" is then

$$R_a^2 = R_a^1 \cap \text{SB}_a(R_b^1).$$

And the event that Ann is rational, Ann strongly believes "Bob is rational," and strongly believes "Bob is rational and strongly believes 'I am rational'" is

$$R_a^3 = R_a \cap \text{SB}_a(R_b) \cap \text{SB}_a(R_b \cap \text{SB}_b(R_a)) = R_a^2 \cap \text{SB}_a(R_b^2).$$

More generally, define R_a^m (resp. R_b^m), so that $R_a^{m+1} = R_a^m \cap \text{SB}_a(R_b^m)$ (resp. $R_b^{m+1} = R_b^m \cap \text{SB}_b(R_a^m)$).

Definition 8. Say there is *rationality and common strong belief of rationality* (RCSBR) at state (s_a, t_a, s_b, t_b) if $(s_a, t_a, s_b, t_b) \in \bigcap_m R_a^m \times \bigcap_m R_b^m$.

The prediction of play under RCSBR is the projection of $\bigcap_m R_a^m \times \bigcap_m R_b^m$ on $S_a \times S_b$. This prediction depends on both the given game and the given epistemic type structure.

Example 3. Return to Example 1, i.e., the BoS with an outside option game and the type structure associated with the lady's choice convention. (Recall, each $\beta_b(t_b)$ assigns probability 1 to $\{Up\} \times \{T_a\}$ and the belief map β_a is onto.) In this example, $\text{proj}_{S_a} R_a^m \times \text{proj}_{S_b} R_b^m$ is $\{Up, Down\} \times \{Out\}$ for each $m \geq 1$.

$m = 1$: Since each type t_b assigns probability 1 to $\{Up\} \times T_a$, (s_b, t_b) is rational if and only if $s_b = Out$. Also, there is a CPS μ_a (resp. ν_a) on $S_b \times T_b$ so that *Up* (resp. *Down*) is sequentially optimal under μ_a (resp. ν_a). Since β_a is onto, there is a type t_a (resp. u_a) so that $(Up, t_a) \in R_a^1$ (resp. $(Down, u_a) \in R_a^1$).

$m \geq 2$: Assume the claim holds for m. Then $R_b^{m+1} \subseteq R_b^m \subseteq \{Out\} \times T_b$. (The second inclusion follows from the induction hypothesis.) Since $R_a^m \cap (\{Up\} \times T_a) \neq \varnothing$, there is a type t_b that assigns probability 1 to R_a^m at each information set. Any such type assigns probability 1 to each R_a^n, for $n \leq m$, at each information set. So $R_b^{m+1} \neq \varnothing$. Thus, $\text{proj}_{S_b} R_b^m = \{Out\}$.

Next, for each $n \leq m$, $\varnothing \neq R_b^n \subseteq \{Out\} \times T_b$. So there is a CPS μ_a with $\mu_a(R_b^m | S_b \times T_b) = 1$. Any such CPS μ_a strongly believes each R_b^n where $n \leq m$. (Here we use the fact that, for each $n \leq m$, $R_b^m \cap (\{In\text{-}Left, In\text{-}Right\} \times T_b) = \varnothing$.) For any such CPS, viz. μ_a, there is a type t_a whose belief is μ_a. As such, there is a type t_a so that $(Up, t_a) \in R_a^{m+1}$ (resp. $(Down, t_a) \in R_a^{m+1}$). ◊

Example 4. Return to Example 2, i.e., the BoS with an outside option game and a complete type structure. In this case, an RCSBR analysis corresponds to the typical forward induction analysis: The strategy *In-Left* is dominated and so there does not exist a type t_b with $(\textit{In-Left}, t_b)$ rational. But for each $s_b \in \{\textit{Out}, \textit{In-Right}\}$, there is a type t_b with (s_b, t_b) rational. Likewise, for each $s_a \in \{\textit{Up}, \textit{Down}\}$, there is a type t_a with (s_a, t_a) rational. It follows that

$$\text{proj}_{S_a} R_a^1 \times \text{proj}_{S_b} R_b^1 = \{\textit{Up}, \textit{Down}\} \times \{\textit{Out}, \textit{In-Right}\}.$$

Now if t_a strongly believes R_b^1, then t_a must assign probability 1 to $\{\textit{In-Right}\} \times T_b$, conditional on BoS being reached. So $\text{proj}_{S_a} R_a^2 \subseteq \{\textit{Down}\}$. Moreover, since β_a is onto, there is a type t_a that strongly believes R_b^1, so

$$\text{proj}_{S_a} R_a^2 \times \text{proj}_{S_b} R_b^2 = \{\textit{Down}\} \times \{\textit{Out}, \textit{In-Right}\}.$$

With this, if t_b strongly believes R_a^2, then t_b must assign probability 1 to *In-Right*, conditional on *In* being played. So $\text{proj}_{S_b} R_b^3 \subseteq \{\textit{In-Right}\}$. Moreover, since β_b is onto, there is a type t_b that strongly believes R_a^2, so

$$\text{proj}_{S_a} R_a^3 \times \text{proj}_{S_b} R_b^3 = \{\textit{Down}\} \times \{\textit{In-Right}\}.$$

A standard induction argument shows that, for each $m \geq 3$, $\text{proj}_{S_a} R_a^m \times \text{proj}_{S_b} R_b^m = \{\textit{Down}\} \times \{\textit{In-Right}\}$. This is the extensive-form rationalizable set. ◊

Comparing Examples 3 and 4 we see that there is a nonmonotonicity in behavioral prediction of RCSBR: even if a type structure contains "more" beliefs, the RCSBR analysis in this "larger" structure can exclude an outcome allowed by an RCSBR analysis in the "smaller" one. To review why this can happen, observe that in the complete type structure (Example 4), there are types of Ann that assign positive probability to Bob's playing *In-Left*, conditional on Ann's information set being reached. But unlike the case of the lady's choice convention (Example 3), no such type can strongly believe the event that Bob is rational. The reason is that, unlike the case of the lady's choice convention, here there are types t_b so that $(\textit{In-Right}, t_b)$ is rational. Thus, in a sense, the nonmonotonicity in the behavioral prediction can be seen as arising from the nonmonotonicity of strong belief.

Example 5. For a given game and epistemic type structure, it may well be the case that $\bigcap_m R_a^m = \varnothing$ and $\bigcap_m R_b^m = \varnothing$. For instance, consider BoS with the outside option and a type structure where $\beta_a(t_a)(\{\textit{In-Left}\} \times T_b | S_b \times T_b) = 1$ for each t_a. Each type of Ann initially assigns positive probability to a strictly dominated strategy of Bob. So $\text{SB}_a(R_b^1) = \varnothing$. Hence, $R_a^2 = \varnothing$. It follows that $\text{SB}_b(R_a^2) = \varnothing$ and so $R_b^3 = \varnothing$. ◊

5. Characterization theorem: EFBRS's

We now turn to characterizing RCSBR. For this it is useful to introduce a *best reply correspondence*, viz. $\rho_a : \mathcal{C}(S_b) \to 2^{S_a}$, where $\rho_a(\mu_a)$ is the set of strategies that are sequentially optimal under μ_a. We begin with extensive-form best response sets.

Definition 9. Call $Q_a \times Q_b \subseteq S_a \times S_b$ an *extensive-form best response set* (EFBRS) if the following hold:

a. For each $s_a \in Q_a$, there is a CPS $\mu_a \in \mathcal{C}(S_b)$ so that

(i) $s_a \in \rho_a(\mu_a)$,

(ii) μ_a strongly believes Q_b, and

(iii) $\rho_a(\mu_a) \subseteq Q_a$.

b. And, likewise, for each $s_b \in Q_b$.

Example 6. Return to BoS with the outside option as in Figure 1. There are three EFBRS: $\{Up, Down\} \times \{Out\}$, $\{Up\} \times \{Out\}$, and $\{Down\} \times \{In\text{-}Right\}$. The first of these is the set of strategies consistent with RCSBR when we append to the game the type structure associated with the lady's choice convention. (See Example 3.) The latter of these is the set of strategies consistent with RCSBR when we append to the game a complete type structure. (See Example 4.) ◊

Why is the EFBRS definition "right" for characterizing RCSBR? Fix some $(s_a, t_a) \in \bigcap R_a^m$. We can immediately identify the first two properties of Definition 9. For the first, recall that s_a is optimal under the CPS associated with t_a, namely $\beta_a(t_a)$. It follows that s_a is optimal under the marginal of $\beta_a(t_a)$ on S_b (a CPS on Bob's strategies). For the second, recall that t_a strongly believes the events R_b^1, R_b^2, R_b^3, etc. So, by the conjunction property of strong belief, t_a strongly believes the event $\bigcap R_b^m$. It then follows from a marginalization property of strong belief that the marginal of $\beta_a(t_a)$ on S_b strongly believes Q_b (i.e., the projection of $\bigcap R_b^m$ onto S_b). Thus, $Q_a \times Q_b$ satisfies both conditions (i) and (ii) of an EFBRS for (s_a, μ_a), where we take μ_a to be the marginal of $\beta_a(t_a)$ on S_b.

But conditions (i) and (ii) do not suffice to characterize RCSBR: We can have a set $Q_a \times Q_b$ that satisfies conditions (i) and (ii) but is inconsistent with RCSBR (for every type structure). This is illustrated by the next example.

Example 7. Consider the game in Figure 2 and the set $Q_a \times Q_b = \{Out\} \times \{Left, Center\}$. We see that the set $Q_a \times Q_b$ satisfies conditions (i) and (ii) of Definition 9. But for each type structure, $\mathrm{proj}_{S_a} \bigcap_m R_a^m \cap \{Out\} = \varnothing$. That is, for each type structure, *Out* is inconsistent with RCSBR.

First we show that $Q_a \times Q_b$ satisfies conditions (i) and (ii) of Definition 9. Begin with Ann and consider the CPS that assigns probability $\frac{1}{2} : \frac{1}{2}$ to *Left* : *Center* at each information set. The strategy *Out* is sequentially optimal under this CPS. Of course, this CPS strongly believes Q_b. Turning to Bob, consider a CPS that assigns probability 1 to *Out* at the initial node and probability $\frac{1}{4} : \frac{1}{4} : \frac{1}{2}$ to *In-Up* : *In-Middle* : *In-Down* conditional on Bob's subgame being reached. The strategies *Left* and *Center* are sequentially optimal under this CPS, and this CPS strongly believes Q_a. So conditions (i) and (ii) are satisfied for $Q_a \times Q_b$.

Next we show that for each type structure, $\mathrm{proj}_{S_a} \bigcap_m R_a^m \cap \{Out\} = \varnothing$. Suppose, contra hypothesis, that there exist some type structure and some type t_a so that $(Out, t_a) \in \bigcap_m R_a^m$. Certainly, (Out, t_a) is rational and t_a strongly believes each R_b^m. Since each

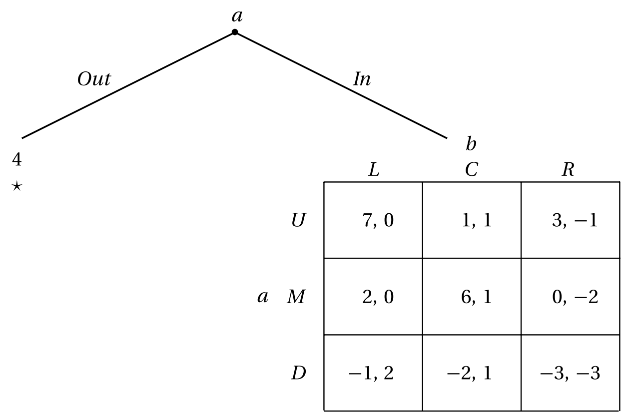

Figure 2. The need for maximality.

pair in $\{Right\} \times T_b$ is irrational and t_a strongly believes "Bob is rational," the type t_a is associated with a CPS that (at each node) assigns probability 1 to $\{Left, Center\} \times T_b$. Now, since (Out, t_a) is rational, the CPS associated with t_a must assign probability $\frac{1}{2} : \frac{1}{2}$ to $\{Left\} \times T_b : \{Center\} \times T_b$ at each node. With this, $(In\text{-}Up, t_a)$ and $(In\text{-}Middle, t_a)$ are also rational. Indeed, since t_a strongly believes each of the R_b^m sets, both $(In\text{-}Up, t_a)$ and $(In\text{-}Middle, t_a)$ must be contained in $\bigcap_m R_a^m$. Now, consider some $(s_b, t_b) \in \bigcap_m R_b^m$. Conditional on Bob's information set being reached, t_b must assign probability 1 to $\{In\text{-}Up, In\text{-}Middle\} \times T_a$. (To see this, note that this event contains rational strategy-type pairs, while the event $\{In\text{-}Down\} \times T_a$ does not contain any rational strategy-type pairs.) Since (s_b, t_b) is rational, $s_b = Center$. Thus, $\bigcap_m R_b^m \subseteq \{Center\} \times T_b$. But, now notice that the CPS associated with t_a does not strongly believe the event $\bigcap_m R_b^m$. By the conjunction property of strong belief, this implies that t_a does not strongly believe some R_m^b, a contradiction. ♢

What went wrong in this example? We began with a set $Q_a \times Q_b$ satisfying conditions (i) and (ii). In particular, we had a strategy $s_a \in Q_a$ for which there was a unique CPS $\mu_a(s_a)$, so that s_a and $\mu_a(s_a)$ satisfy conditions (i) and (ii). But there was also a strategy $r_a \in S_a \setminus Q_a$ that was sequentially optimal under $\mu_a(s_a)$. (Actually, there were two such strategies.) As a result, if (s_a, t_a) is consistent with RCSBR, then (r_a, t_a) must also be consistent with RCSBR. Thus, there may be a strategy of Ann that is consistent with RCSBR, but is not contained in Q_a. And, if so, we may be able to find an s_b and a CPS $\mu_b(s_b)$ (on S_a) so that s_b and $\mu_b(s_b)$ satisfy conditions (i) and (ii), despite the fact that s_b is not optimal under any CPS (on $S_a \times T_a$) that strongly believes the RCSBR strategy-type pairs for Ann.

This suggests that we need to add a maximality criterion to conditions (i) and (ii) of Definition 9. Indeed, this is what condition (iii) achieves.

Theorem 1. *Fix an extensive-form game* Γ.

(i) For any Γ*-based type structure,* $\text{proj}_{S_a} \bigcap_m R_a^m \times \text{proj}_{S_b} \bigcap_m R_b^m$ *is an EFBRS.*

(ii) Fix a nonempty EFBRS $Q_a \times Q_b$. There exists a Γ-based type structure, so that $Q_a \times Q_b = \mathrm{proj}_{S_a} \bigcap_m R_a^m \times \mathrm{proj}_{S_b} \bigcap_m R_b^m$.

Proof. Begin by showing part (i) of the theorem. Fix a Γ-based type structure. If $\bigcap_m R_a^m \times \bigcap_m R_b^m = \varnothing$, then the result is immediate. So suppose $\bigcap_m R_a^m \times \bigcap_m R_b^m \neq \varnothing$. Fix $(s_a, s_b) \in \mathrm{proj}_{S_a} \bigcap_m R_a^m \times \mathrm{proj}_{S_b} \bigcap_m R_b^m$. Then there exists (t_a, t_b) such that

$$(s_a, t_a, s_b, t_b) \in \bigcap_m R_a^m \times \bigcap_m R_b^m.$$

We show that the CPS $\mathrm{marg}_{S_b} \beta_a(t_a)$ satisfies conditions (i)–(iii) of an EFBRS for the strategy s_a. A similar argument holds for s_b.

Begin with the fact that

$$(s_a, t_a) \in \rho_a(\mathrm{marg}_{S_b} \beta_a(t_a)) \times \{t_a\} \subseteq R_a.$$

Now use the fact that t_a strongly believes each R_b^m to get that

$$\rho_a(\mathrm{marg}_{S_b} \beta_a(t_a)) \times \{t_a\} \subseteq \bigcap_m R_a^m.$$

So, $s_a \in \rho_a(\mathrm{marg}_{S_b} \beta_a(t_a)) \subseteq \mathrm{proj}_{S_a} \bigcap_m R_a^m$, establishing conditions (i) and (iii) of an EFBRS. Next, use the conjunction property of strong belief (Property 1) to get that $\beta_a(t_a)$ strongly believes $\bigcap_m R_b^m$. Using the marginalization property (Property 2), $\mathrm{marg}_{S_a} \beta_a(t_a)$ strongly believes $\mathrm{proj}_{S_b} \bigcap_m R_b^m$. This establishes condition (ii) of an EFBRS.

Now turn to part (ii) of the theorem. Fix an EFBRS $Q_a \times Q_b \neq \varnothing$. Let $T_a = Q_a$ and $T_b = Q_b$. Fix a type $t_a \in T_a = Q_a$. There is a CPS $\mu_a(t_a) \in \mathcal{C}(S_b)$ satisfying conditions (i)–(iii) of an EFBRS. Now construct a CPS $\beta_a(t_a) \in \mathcal{C}(S_b \times T_b, \mathcal{S}_b \otimes T_b)$ as follows. If $Q_b \cap S_b(h) \neq \varnothing$, set $\beta_a(t_a)((t_b, t_b)|S_b(h) \times T_b) = \mu_a(t_a)(t_b|S_b(h))$ for each $t_b \in Q_b = T_b$. Next fix some arbitrary element $t_b^* \in T_b$. If $Q_b \cap S_b(h) = \varnothing$, set $\beta_a(t_a)((s_b, t_b^*)|S_b(h) \times T_b) = \mu_a(t_a)(s_b|S_b(h))$ for each $s_b \in S_b$. (Type t_b^* is the same for each information set with $Q_b \cap S_b(h) = \varnothing$.)

Indeed, each $\beta_a(t_a)$ is a CPS on $\mathcal{S}_b \otimes T_b$. Conditions (i) and (ii) of a CPS are immediate. For condition (iii), fix an event E_b and two information sets $h, i \in H_a$ with $E_b \subseteq S_b(h) \times T_b \subseteq S_b(i) \times T_b$. First, consider the case where $Q_b \cap S_b(h) \neq \varnothing$. In this case, $Q_b \cap S_b(i) \neq \varnothing$. So

$$\begin{aligned}
\beta_a(t_a)(E_b|S_b(i) \times T_b) &= \mu_a(t_a)\big(\{t_b \in Q_b : (t_b, t_b) \in E_b\}|S_b(i)\big) \\
&= \mu_a(t_a)\big(\{t_b \in Q_b : (t_b, t_b) \in E_b\}|S_b(h)\big) \times \mu_a(t_a)(S_b(h)|S_b(i)) \\
&= \mu_a(t_a)\big(\{t_b \in Q_b : (t_b, t_b) \in E_b\}|S_b(h)\big) \times \mu_a(t_a)(Q_b \cap S_b(h)|S_b(i)) \\
&= \beta_a(t_a)(E_b|S_b(h) \times T_b) \times \beta_a(t_a)(S_b(h) \times T_b|S_b(i) \times T_b),
\end{aligned}$$

where the first and fourth lines follow from the construction, the second line follows from the fact that $\mu_a(t_a)$ is a CPS, and the third line follows from the fact that $\mu_a(t_a)(Q_b|S_b(h)) = 1$ (since $Q_b \cap S_b(h) \neq \varnothing$ and $\mu_a(t_a)$ strongly believes Q_b). This establishes condition (iii) of a CPS when $Q_b \cap S_b(h) \neq \varnothing$. So suppose $Q_b \cap S_b(h) = \varnothing$

and recall $E_b \subseteq S_b(h) \times T_b$. If $Q_b \cap S_b(i) \neq \emptyset$, then $\mu_a(t_a)(\text{proj}_{S_b} E_b|S_b(i)) = 0$ and $\mu_a(t_a)(S_b(h)|S_b(i)) = 0$. (This uses the fact that $\mu_a(t_a)(Q_b|S_b(i)) = 1$, which follows from strong belief.) So, here too,

$$\begin{aligned}\beta_a(t_a)(E_b|S_b(i) \times T_b) &= \beta_a(t_a)(E_b|S_b(h) \times T_b) \times \beta_a(t_a)(S_b(h) \times T_b|S_b(i) \times T_b)\\ &= 0.\end{aligned}$$

Finally, suppose $Q_b \cap S_b(i) = \emptyset$. Here

$$\begin{aligned}\beta_a(t_a)(E_b|S_b(i) \times T_b) &= \mu_a(t_a)\big(\{s_b : (s_b, t_b^*) \in E_b\}|S_b(i)\big)\\ &= \mu_a(t_a)\big(\{s_b : (s_b, t_b^*) \in E_b\}|S_b(h)\big) \times \mu_a(t_a)(S_b(h)|S_b(i))\\ &= \beta_a(t_a)(E_b|S_b(h) \times T_b) \times \beta_a(t_a)(S_b(h) \times \{t_b^*\}|S_b(i) \times T_b)\\ &= \beta_a(t_a)(E_b|S_b(h) \times T_b) \times \beta_a(t_a)(S_b(h) \times T_b|S_b(i) \times T_b),\end{aligned}$$

as required.

We conclude the proof by showing

$$Q_a = \bigcup_{t_a \in T_a} [\rho_a(\text{marg}_{S_b} \beta_a(t_a))] \tag{1}$$

$$R_a^m = \bigcup_{t_a \in T_a} \big[\rho_a(\text{marg}_{S_b} \beta_a(t_a)) \times \{t_a\}\big] \quad \text{for each } m, \tag{2}$$

and likewise with a and b interchanged. Taken together, they give the desired result.

Part (1): Recall that for each $t_a \in T_a = Q_a$, $\mu_a(t_a) = \text{marg}_{S_b} \beta_a(t_a)$. So it is immediate from the construction that $Q_a \subseteq \bigcup_{t_a \in T_a} \rho_a(\text{marg}_{S_b} \beta_a(t_a))$. Conversely, fix any strategy s_a in $\bigcup_{t_a \in T_a} \rho_a(\text{marg}_{S_b} \beta_a(t_a))$. Then there is a type $t_a \in T_a = Q_a$ so that s_a is sequentially optimal under $\mu_a(t_a)(\cdot|\cdot)$. It follows from part (iii) of the definition of an EFBRS that $s_a \in Q_a$.

Part (2): The proof is by induction on m. The equation is immediate for $m = 1$. Assume the result holds for m. To show that it holds for $m + 1$, it suffices to show that each $t_a \in T_a$ strongly believes R_b^m. For this, fix an information set h such that $R_b^m \cap [S_b(h) \times T_b] \neq \emptyset$. Observe that

$$\begin{aligned}[\text{proj}_{S_b} R_b^m] \cap S_b(h) &= \Big[\bigcup_{t_b \in T_b} \rho_b(\text{marg}_{S_a} \beta_b(t_b))\Big] \cap S_b(h)\\ &= Q_b \cap S_b(h).\end{aligned}$$

(The first equality follows from the induction hypothesis for b; the second equality follows from (1).) Since $R_b^m \cap [S_b(h) \times T_b] \neq \emptyset$, it follows that $Q_b \cap S_b(h) \neq \emptyset$ and so $\mu_a(t_a)(Q_b|S_b(h)) = 1$. (Here, we use part (ii) of the definition of an EFBRS.) So, by construction, $\beta_a(t_a)(R_b^m|S_b(h) \times T_b) = 1$, as required. □

Part (i) of Theorem 1 says that the projection of the RCSBR event on $S_a \times S_b$ is an EFBRS. But this may form an empty EFBRS. That said, there is always a nonempty EFBRS.

Remark 1. For any game, there exists a nonempty EFBRS—namely, the set of extensive-form rationalizable strategy profiles.

Battigalli and Siniscalchi (1999a) show that for each Γ, there exists a complete Γ-based type structure with compact metrizable type sets.[4] Proposition 6 in Battigalli and Siniscalchi (2002) says that for each such complete structure, the projection of the RCSBR event onto $S_a \times S_b$ is the set of extensive-form rationalizable strategies. So using Theorem 1(i), this set is an EFBRS. The fact that it is nonempty is shown as Corollary 1 in Battigalli (1997).

6. Alternate characterization theorem: Directed rationalizability

Return to the lady's choice convention example, i.e., Example 1. There, each type of Bob is associated with some CPS that assigned probability 1 to $\{Up\} \times T_a$. This gives a restriction on Bob's first-order beliefs, i.e., his beliefs about what Ann chooses. Let Δ_b represent this restriction on first-order beliefs. So Δ_b is a subset of the CPS's on S_a and, in our example, Δ_b contains only the CPS that assigns probability 1 to *Up* at each information set. We do not have a restriction on Ann's first-order beliefs. So we write Δ_a for the set of all CPS's on S_b.

With $\Delta = \Delta_a \times \Delta_b$ in hand, we can take an iterative approach to analyzing the game tree—much like a "typical rationalizability" procedure. In round one, we eliminate *In-Left* and *In-Right* for Bob, since these strategies are not sequentially optimal under the CPS in Δ_b. We do not eliminate any of Ann's strategies, since they are each sequentially optimal under some CPS (in Δ_a). So in round one, we are left with the set $\{Up, Down\} \times \{Out\}$. Turning to round two, *Out* is sequentially optimal under the CPS in Δ_b and that CPS strongly believes $\{Up, Down\}$. Thus, we cannot eliminate *Out* in round two. Likewise, *Up* (resp. *Down*) is sequentially optimal under a CPS that assigns probability 1 to *Out* at the initial node and probability 1 to *Left* (resp. *Right*) at Bob's subgame. This CPS is contained in Δ_a and strongly believes $\{Out\}$. So we also get $\{Up, Down\} \times \{Out\}$ in round two. Indeed, a standard induction argument gives that $\{Up, Down\} \times \{Out\}$ is the outcome of the procedure. Of course, this is the EFBRS we identify in Section 4.

The procedure we use above is called Δ-*rationalizability*; see Battigalli and Siniscalchi (2003).[5] More generally, let Δ_a (resp. Δ_b) be a nonempty subset of $\mathcal{C}(S_b)$ (resp. $\mathcal{C}(S_a)$), i.e., a set of first-order beliefs of Ann (resp. Bob). Call $\Delta = \Delta_a \times \Delta_b$ a *set of first-order beliefs*. Set $S_a^{\Delta,0} = S_a$ and $S_b^{\Delta,0} = S_b$. Inductively define $S_a^{\Delta,m}$ and $S_b^{\Delta,m}$ as follows: Let $S_a^{\Delta,m+1}$ be the set of all $s_a \in S_a^{\Delta,m}$ so that there is some CPS $\mu_a \in \Delta_a$ with (i) $s_a \in \rho_a(\mu_a)$ and (ii) μ_a strongly believes $S_b^{\Delta,1}, \ldots, S_b^{\Delta,m}$. And likewise with a and b interchanged.[6]

[4]Battigalli and Siniscalchi (1999a) canonical construction is a type structure in the sense of Definition 2. Specifically, in the case of a game tree, the basic conditioning events are clopen and so Battigalli and Siniscalchi (1999a) get T_a and T_b to be compact metrizable as an output.

[5]Battigalli and Siniscalchi (2003) use the concept to study a different problem from the one studied here. In their problem, the set Δ is given to the analyst. In our problem, Δ may be unknown to the analyst and we obtain a characterization across all Δ's. See Section 9.a.

[6]This definition is as in Battigalli (1999). It is a stronger requirement than the definition in Battigalli and Siniscalchi (2003). They put $s_a \in S_a^{\Delta,m+1}$ if $s_a \in S_a^{\Delta,m}$ and there is some CPS $\mu_a \in \Delta_a$ with (i) $s_a \in \rho_a(\mu_a)$ and

Definition 10 (Battigalli and Siniscalchi 2003). Call $S_a^{\Delta} = \bigcap_{m\geq 0} S_a^{\Delta,m}$ (resp. $S_b^{\Delta} = \bigcap_{m\geq 0} S_b^{\Delta,m}$) the Δ-*rationalizable strategies of Ann* (resp. *Bob*). Call $S_a^{\Delta} \times S_b^{\Delta}$ the Δ-*rationalizable strategy set*.

Since the sets $S_a^{\Delta,m} \times S_b^{\Delta,m}$ form a decreasing sequence and $S_a \times S_b$ is finite, there is some (finite) M so that $S_a^{\Delta} \times S_b^{\Delta} = S_a^{\Delta,M} \times S_b^{\Delta,M}$.

Of course, there may be many Δ-rationalizable sets, each of which is obtained by beginning the procedure with a different set of first-order beliefs $\Delta = \Delta_a \times \Delta_b$. We use the phrase *directed rationalizability* to refer to the set of all $S_a^{\Delta} \times S_b^{\Delta}$. So, for a given game Γ, the directed rationalizability concept gives $\{S_a^{\Delta} \times S_b^{\Delta} : \Delta = \Delta_a \times \Delta_b \subseteq \mathcal{C}(S_b) \times \mathcal{C}(S_b)\}$.

Beginning from the lady's choice example, we can use the type structure to construct an associated set of first-order beliefs Δ and this set of first-order beliefs Δ can be used to perform Δ-rationalizability. The output is the EFBRS we identified earlier. But the lady's choice convention has a particular feature: it is a restriction on first-order beliefs and a requirement that the restriction be "transparent" to the players. So the only restriction on second-order beliefs (i.e., beliefs about strategy the other player chooses and the other player's the first-order beliefs) is the requirement that at each information set, Ann must believe that Bob believes she will play *Up* and so on. It is this transparency of (only) first-order restrictions that allows us to directly compute the associated directed rationalizability set.

More generally, when we begin from a given type structure, we impose substantive assumptions about which beliefs players do versus do not consider possible. These assumptions may correspond to restrictions (only) on players' first-order beliefs, which are transparent to the players. But they need not: they may involve additional restrictions on higher-order beliefs, and if they do, the procedure we outline above fails.

To see the failure, begin with an epistemic type structure and use the structure itself to form the set $\bar{\Delta} = \bar{\Delta}_a \times \bar{\Delta}_b$. Specifically, for each type $t_a \in T_a$, consider the marginal of $\beta_a(t_a)$ on S_b. These CPS's form the set $\bar{\Delta}_a$. Construct the set $\bar{\Delta}_b$ analogously. Here, the strategies that survive one round of $\bar{\Delta}$-rationalizability are exactly the strategies that are consistent with $\text{R0SBR}_a \times \text{R0SBR}_b$. But, in round two, we lose the equivalence: If $\beta_a(t_a)$ strongly believes the event "Bob is rational," then the marginal of $\beta_a(t_a)$ also strongly believes that "Bob chooses a strategy consistent with one round of elimination of $\bar{\Delta}$-rationalizability." (Here, we use a marginalization property of strong belief, plus the round-one equivalence.) But the converse need not hold. So the strategies that survive two rounds of $\bar{\Delta}$-rationalizability may strictly contain the R1SBR strategies. And on round three, we loose the inclusion. If the CPS$\beta_a(t_a)$ strongly believes the R1SBR event for Bob, then the marginal of $\beta_a(t_a)$ also strongly believes that "Bob chooses a strategy consistent with R1SBR." But recall that the strategies consistent with R1SBR may

(ii) μ_a strongly believes $S_b^{\Delta,m}$. Any set that satisfies the requirements here also satisfies the requirements in Battigalli and Siniscalchi (2003), but the converse does not hold. (See Battigalli and Prestipino 2011 for an example.) Thus, using Theorem 1 here, it can be shown that the definition of Battigalli and Siniscalchi (2003) is conceptually incorrect. (Battigalli and Prestipino 2011 point out that the two definitions are equivalent when Δ satisfies a "closedness under composition" condition. Since Battigalli and Siniscalchi 2003 focus on the case where this condition is satisfied, their results hold with the definition given here.)

be strictly contained in the strategies that survive two rounds of $\bar{\Delta}$-rationalizability. So there may be information sets consistent with this latter event, but not the former. This implies that even if $\beta_a(t_a)$ strongly believes the R1SBR event for Bob, it need not strongly believe that Ann's behavior is consistent with two rounds of $\bar{\Delta}$-rationalizability. (This is an instance of the fact that strong belief is not monotonic.) As such, we can lose (any) relationship between the RCSBR strategies and the $\bar{\Delta}$-rationalizable strategy set. In fact, Appendix B illustrates an example where the RCSBR strategy set and the $\bar{\Delta}$-rationalizable strategy set are disjoint.

There is another route that instead uses the EFBRS properties to form a set $\Delta = \Delta_a \times \Delta_b$ of first-order beliefs. Fix an epistemic structure. The RCSBR strategies form an EFBRS, viz. $Q_a \times Q_b$. For each $s_a \in Q_a$, we have some CPS $\mu_a(s_a)$ satisfying the conditions of an EFBRS. Take Δ_a to be the set of such CPS's, i.e., one for each $s_a \in Q_a$, and construct Δ_b similarly. Now we do have an equivalence between the RCSBR strategies and the Δ-rationalizable strategies. More precisely, for each $m \geq 1$, $Q_a \times Q_b$ is the set of strategies that survives m-rounds of elimination of Δ-rationalizability. The case of $m = 1$ follows from properties (i) and (iii) of an EFBRS, the case of $m = 2$ uses condition (ii) of an EFBRS, and so on, by induction.

Proposition 1. *Fix an extensive-form game* Γ.

(i) Given an EFBRS, viz. $Q_a \times Q_b$, *there exists a set of first-order beliefs, viz.* $\Delta = \Delta_a \times \Delta_b$, *so that* $S_a^\Delta \times S_b^\Delta = Q_a \times Q_b$.

(ii) Given a set of first-order beliefs, viz. $\Delta = \Delta_a \times \Delta_b$, $S_a^\Delta \times S_b^\Delta$ *is an EFBRS.*

Thus, in conjunction with Theorem 1, we have the following alternate characterization theorem.

Corollary 1. *Fix an extensive-form game* Γ.

(i) For any Γ*-based type structure, there exists a set of first-order beliefs, viz.* $\Delta = \Delta_a \times \Delta_b$, *so that* $S_a^\Delta \times S_b^\Delta = \text{proj}_{S_a} \bigcap_m R_a^m \times \text{proj}_{S_b} \bigcap_m R_b^m$.

(ii) Fix a set of first-order beliefs, viz. $\Delta_a \times \Delta_b$. *Then there exists a* Γ*-based structure so that* $S_a^\Delta \times S_b^\Delta = \text{proj}_{S_a} \bigcap_m R_a^m \times \text{proj}_{S_b} \bigcap_m R_b^m$.

Proof of Proposition 1. Begin with part (i). Fix an EFBRS set $Q_a \times Q_b$. For each $s_a \in Q_a$, there exists a corresponding CPS $\mu_a(s_a) \in \mathcal{C}(S_b)$ satisfying conditions (i)–(iii) of an EFBRS for $Q_a \times Q_b$. Take Δ_a so that, for each $s_a \in Q_a$, Δ_a contains exactly one such CPS $\mu_a(s_a)$. There are no other CPS's in Δ_a. Define Δ_b analogously. We show that for each $m \geq 1$, $S_a^{\Delta,m} \times S_b^{\Delta,m} = Q_a \times Q_b$. This establishes the result.

The proof is by induction. Begin with $m = 1$. Certainly $Q_a \subseteq S_a^{\Delta,1}$. Fix $s_a \in S_a^{\Delta,1}$. Then there exists some $\mu_a \in \Delta_a$ so that s_a is sequentially optimal under μ_a. This CPS μ_a is associated with some $r_a \in Q_a$, i.e., so that r_a and μ_a jointly satisfy conditions (i)–(iii) of an EFBRS. Now apply condition (iii) of an EFBRS to get that $s_a \in Q_a$.

Now fix $m \geq 2$ and assume $S_a^{\Delta,n} \times S_b^{\Delta,n} = Q_a \times Q_b$ for all $n \leq m$. We show that it also holds for $m+1$. Fix $s_a \in Q_a = S_a^{\Delta,m}$. Then using the construction of Δ_a, there exists some $\mu_a \in \Delta_a$ satisfying conditions (i) and (ii) of an EFBRS for $Q_a \times Q_b$, so that $s_a \in \rho_a(\mu_a)$ and μ_a strongly believes $Q_b = S_b^{\Delta,n}$ for all $n \leq m$. So certainly, $Q_a \subseteq S_a^{\Delta,m+1}$. Conversely, fix some $s_a \in S_a^{\Delta,m+1}$. Then there exists a CPS $\mu_a \in \Delta_a$ so that $s_a \in \rho_a(\mu_a)$ and μ_a strongly believes $S_b^{\Delta,m}$. Again, since each element of Δ_a satisfies conditions (i)–(iii) of an EFBRS for some $r_a \in Q_a$, it follows that $\rho_a(\mu_a) \subseteq Q_a$ and so $s_a \in Q_a$.

Now turn to part (ii) of the proposition. Fix some set of first-order beliefs, viz. $\Delta = \Delta_a \times \Delta_b$. There exist some M with $S_a^\Delta \times S_b^\Delta = S_a^{\Delta,M} \times S_b^{\Delta,M}$. Fix $s_a \in S_a^\Delta$. There exists a CPS μ_a so that $s_a \in \rho_a(\mu_a)$ and μ_a strongly believes each $S_b^{\Delta,m}$ for $m \leq M$. Thus s_a and μ_a satisfy conditions (i) and (ii) of an EFBRS for $Q_a \times Q_b = S_a^\Delta \times S_b^\Delta$. Moreover, if $r_a \in \rho_a(\mu_a)$, then r_a is optimal under a CPS that strongly believes each $S_b^{\Delta,m}$ for $m \leq M$. As such, $r_a \in S_a^{\Delta,m}$ for each $m \leq M$, establishing that $r_a \in S_a^\Delta$. Therefore, condition (iii) of an EFBRS is also satisfied. A similar argument applies to b. Therefore, $S_a^\Delta \times S_b^\Delta$ is an EFBRS. □

The proof of Proposition 1 gives an ancillary result. Begin with some finite set of first-order beliefs, viz. $\Delta = \Delta_a \times \Delta_b$. Proposition 1(ii) says that $S_a^\Delta \times S_b^\Delta$ is an EFBRS. Conversely, begin with some EFBRS. The proof of Proposition 1(i) says that we can find a finite set of first-order beliefs, viz. $\Delta = \Delta_a \times \Delta_b$, so that $S_a^\Delta \times S_b^\Delta$ is this EFBRS.

Remark 2. Fix a game tree Γ. The directed rationalizability set is

$$\{S_a^\Delta \times S_b^\Delta : \Delta = \Delta_a \times \Delta_b \subseteq \mathcal{C}(S_b) \times \mathcal{C}(S_b)\} = \{S_a^\Delta \times S_b^\Delta : \Delta = \Delta_a \times \Delta_b \text{ is finite}\}.$$

Thus, using the EFBRS properties, we can see that we need only to compute the Δ-rationalizable sets for finite sets of first-order beliefs. Of course, much as is the case with EFBRS's, the Δ-rationalizable strategy set may be empty. When $\Delta = \mathcal{C}(S_a) \times \mathcal{C}(S_b)$, $S_a^\Delta \times S_b^\Delta$ is the extensive-form rationalizable strategy set. So in keeping with Remark 1, there always exists a nonempty Δ-rationalizable strategy set.

While the EFBRS and directed rationalizability concepts are equivalent, it often is useful to focus on the former definition. The reason is that properties (i), (ii), and (iii) of an EFBRS give some immediate implications in terms of behavior. In Sections 7 and 8, we discuss the consequences of context-dependent forward reasoning for some specific games. There the EFBRS properties play an important role, much in the same way that the properties of a self-admissible set (Brandenburger et al. 2008) play an important role in analyzing games. Indeed, we see that these properties help to analyze games such as centipede, the finitely repeated prisoner's dilemma, and perfect information games.

7. Analyzing games

In this section, we analyze the predictions of RCSBR in games of interest. We do so by making use of the properties of an EFBRS and not the (equivalent) directed rationalizability definition.

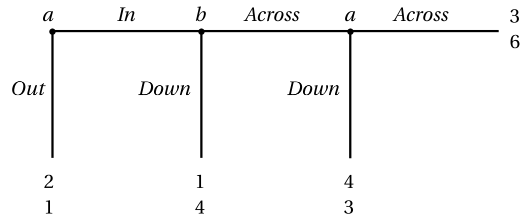

FIGURE 3. Three-legged centipede.

a \ b	C	D
C	c, c	e, d
D	d, e	$0, 0$

$d > c > 0 > e$

FIGURE 4. Prisoner's dilemma.

EXAMPLE 8. Consider the three-legged centipede game given in Figure 3. Here, the EFBRS's are $\{Out\} \times \{Down\}$ and $\{Out\} \times \{Down, Across\}$. In particular, there is no EFBRS where Ann plays *In* at the first node. To see this, suppose otherwise, i.e., suppose there exists an EFBRS $Q_a \times Q_b$ and a strategy $s_a \in Q_a$, where s_a plays *In* at the first node. By condition (i) of an EFBRS, we must have that $Q_a \subseteq \{Out, In\text{-}Down\}$, so that $s_a = In\text{-}Down$. Now, fix $s_b \in Q_b$ and recall that s_b must be sequentially optimal under a CPS that strongly believes Q_a. Then, at Bob's information set, this CPS must assign probability 1 to *In-Down*. Since s_b is sequentially optimal under this CPS, $s_b = Down$. So we have that $Q_b = \{Down\}$. But then *In-Down* cannot simultaneously satisfy conditions (i) and (ii) of an EFBRS. ◊

The argument we present for the three-legged centipede is more general: Fix an EFBRS of the n-legged centipede game. Then the first player chooses *Out*. This result is a consequence of Proposition 3(i) to come.

EXAMPLE 9. Figure 4 gives the prisoner's dilemma. Consider the 3-repeated version of the game. Let $Q_a \times Q_b$ be a nonempty EFBRS. Then each $(s_a, s_b) \in Q_a \times Q_b$ results in the *Defect-Defect* path.[7]

Let us give an intuition: By condition (i) of an EFBRS, each strategy $s_a \in Q_a$ (resp. $s_b \in Q_b$) is sequentially justifiable. As such, s_a (resp. s_b) plays *Defect* in the last period at each history allowed by s_a (resp. s_b). Now consider a second period information set h, where $s_a \in S_a(h)$ and $Q_b \cap S_b(h) \neq \emptyset$. By conditions (i) and (ii) of an EFBRS, s_a must be sequentially optimal under a CPS $\mu_a(s_a)$ with $\mu_a(s_a)(Q_b | S_b(h)) = 1$. Then, conditional

[7] In the once or twice repeated prisoner's dilemma, we have a stronger result: If (s_a, s_b) is contained in an EFBRS, then each of s_a and s_b specify *Defect* at each information set.

on h, $\mu_a(s_a)$ assigns probability 1 to Bob defecting in the third period, irrespective of Ann's play. As such, s_a plays D at h. And likewise with a and b reversed.

Turn to the first period and suppose, contra hypothesis, there is some $s_a \in Q_a$ so that s_a initially chooses C. For each $s_a \in Q_b$, (s_a, s_b) results in the *Defect-Defect* path in periods two and three. So Ann's expected payoffs from s_a corresponds to her first period expected payoffs from playing s_a. With this, the *Defect*-always strategy yields a strictly higher expected payoff in the first period and an expected payoff of at least zero in subsequent periods. This contradicts s_a being optimal under $\mu_a(s_a)(\cdot|S_b)$. $\Diamond$

An analogous result holds for the N-repeated prisoner's dilemma for N finite. The proof is given in Appendix C.

Let us take stock of the examples above. First, in battle of the sexes with the outside option, we get that either (i) Bob plays *Out* or (ii) Bob plays *In-Right* and Ann plays *Down*. Each of these were subgame perfect paths of play. In centipede, we get the backward induction path (but not necessarily the backward induction strategies). Likewise, in the finitely repeated prisoner's dilemma, we get the unique Nash (and so subgame perfect) path, where each player plays *Defect* in all periods.

In each of these cases, the outcomes allowed by an EFBRS coincide with the outcomes allowed by some subgame perfect equilibrium (SPE). This raises the question, Are the EFBRS concept and the SPE concept equivalent? If so, then we have a good idea what the EFBRS concept delivers (in games of interest), since we have a good idea about what SPE delivers.

The EFBRS and SPE concepts are not equivalent, but in a particular class of games, any pure-strategy SPE corresponds to some EFBRS. Each of the examples we mentioned is contained in this class of games.

Definition 11. Say a game Γ has *observable actions* if each information set is a singleton.

To understand the definition, recall that in our setup, both a and b have a choice at each history. (Of course, it may be the case that only one of the players is active.) So a game with observable actions is a game where the players begin by making simultaneous choices, learn the realization of the choices, and then perhaps make simultaneous choices, etc., until a terminal history is reached.

Given distinct terminal histories, viz. z and z', we can write $z = (x, c^1, \ldots, c^K)$ and $z' = (x, d^1, \ldots, d^L)$, where x is the last common predecessor of z and z', i.e., $c^1 \neq d^1$. (Recall, $c^k = (c_a^k, c_b^k)$ and $d^l = (d_a^l, d_b^l)$.)

Definition 12. Fix a game of observable actions and two distinct terminal nodes, viz. $z = (x, c^1, \ldots, c^K)$ and $z' = (x, d^1, \ldots, d^L)$. Say *$a$ is decisive for* (z, z') if a moves at x, $c_a^1 \neq d_a^1$, and $c_b^1 = d_b^1$. And likewise with a and b interchanged.

Definition 13 (Battigalli 1997). A game of observable actions satisfies *no relevant ties* (NRT) if, whenever a (resp. b) is decisive for (z, z'), then $\Pi_a(z) \neq \Pi_a(z')$.

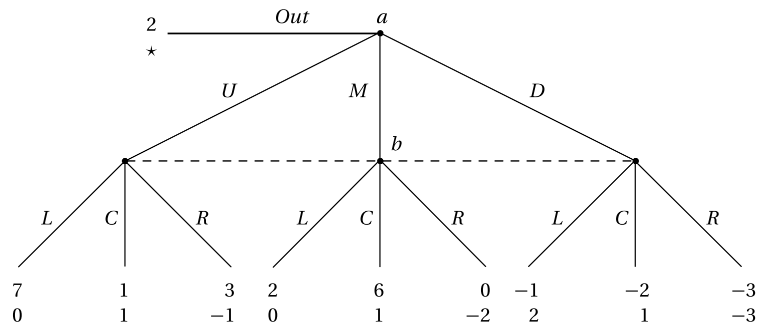

Figure 5. A modification of Figure 2.

A game with no ties satisfies NRT, but the converse does not hold. Reny's (1993, Figure 1) take-it-or-leave-it game is one such example.

Fix a strategy s_a and write $[s_a]$ for the set of all r_a that induce the same plan of action as s_a, i.e., the set of all r_a so that $\zeta(r_a, \cdot) = \zeta(s_a, \cdot)$, and likewise define $[s_b]$.

Proposition 2. *Fix a game* Γ *with observable actions and a pure-strategy SPE, viz.* (s_a, s_b).

(i) There is an EFBRS, viz. $Q_a \times Q_b$*, so that* $[s_a] \times [s_b] \subseteq Q_a \times Q_b$.

(ii) If Γ *satisfies NRT, then* $[s_a] \times [s_b]$ *is an EFBRS.*

Each of the examples we have seen satisfies both observable actions and NRT. In those examples, any pure-strategy subgame perfect equilibrium (s_a, s_b) belongs to an EFBRS, where the EFBRS only allows the terminal node $\zeta(s_a, s_b)$. This fits with part (ii) of the proposition. Part (i) says that even if the game fails NRT, (s_a, s_b) still is contained in some EFBRS. Example 12 in Appendix C provides a game that fails NRT, so that any EFBRS that contains a certain pure-strategy SPE also allows other paths of play.

Proposition 2 does not say that the pure-strategy SPE concept and the EFBRS concept are equivalent. A game without observable actions may have a pure-strategy subgame perfect equilibrium whose outcome is precluded by any EFBRS. Conversely, a given EFBRS may allow outcomes that are precluded by any (even randomized) subgame perfect equilibrium. (This can happen even in a game with observable actions and NRT.) The next examples demonstrate these points.

Example 10. The game in Figure 5 satisfies NRT but fails the observable actions condition. It is obtained from the game in Figure 2 by two transformations. First, the simultaneous move subgame is transformed into a game where Ann moves first and then Bob moves not knowing Ann's choice. Second, two of Ann's decision nodes are coalesced.

Here, $(Out, Right)$ is a pure-strategy subgame perfect equilibrium. But applying the argument in Section 5, *Out* is not contained in any EFBRS.[8] ◊

[8] Unlike the subgame perfect concept, the EFBRS concept is invariant to coalescing decision nodes.

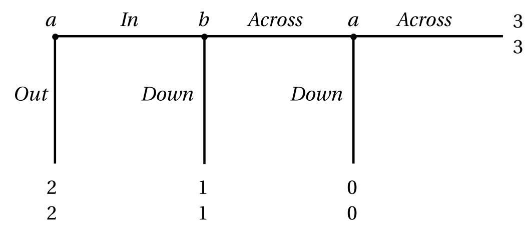

FIGURE 6. A common interest game.

EXAMPLE 11. The game in Figure 6 satisfies both NRT and the observable actions condition. The unique subgame perfect equilibrium is (*In-Across*, *Across*), which results in the $(3, 3)$ outcome. Indeed, this profile induces an EFBRS, viz. $\{$*In-Across*$\} \times \{$*Across*$\}$. But there are two EFBRS's that give the $(2, 2)$ outcome, namely $\{$*Out*$\} \times \{$*Down*$\}$ and $\{$*Out*$\} \times \{$*Down*, *Across*$\}$. ◊

Taken together with the main theorem (Theorem 1), Example 11 says that a non-backward induction outcome, namely $(2, 2)$, is consistent with RCSBR. To understand this better, notice that *Out* is the unique best response for Ann under a CPS that assigns probability 1 to the event "Bob plays *Down*." So if each type of Ann assigns probability 1 to $\{$*Down*$\} \times T_b$, then conditional on Bob's node being reached, he must conclude that Ann is irrational. In this case, Bob may very well believe that Ann is playing *In-Down*; if so, *Down* is a unique (sequential) best response for Bob.

8. PERFECT-INFORMATION GAMES

Example 10 shows that in games without observable actions, the SPE concept allows for outcomes that are excluded by every EFBRS. Alternatively, Proposition 2 and Example 11 show that in games with observable actions, the SPE concept is a strict refinement of the EFBRS concept. Thus, even in these games, we cannot use the SPE concept to analyze the consequences of RCSBR.

Now we turn to a particular class of games with observable actions, namely perfect-information games (i.e., games with observable actions and with at most one active player at each information set). We have seen some examples of perfect-information games, e.g., Examples 8 and 11. In the former case, each EFBRS yields the backward induction path (and so the backward induction outcome). Of course, for that game, the Nash and backward induction paths coincide. Alternatively, in Example 11, one EFBRS corresponds to backward induction, but others do not. However, there we do get that the EFBRS paths correspond (exactly) to the Nash paths (and so Nash outcomes) of the game.

The examples suggest there may be a connection between EFBRS's and Nash outcomes, at least for perfect-information (PI) games. (Of course, for non-PI games, an EFBRS may give non-Nash outcomes.) Indeed, there is a connection for PI games satisfying a "no ties" condition.

Definition 14 (Marx and Swinkels 1997). A game satisfies *transference of decision-maker indifference* (TDI) if $\pi_a(s_a, s_b) = \pi_a(r_a, s_b)$ implies $\pi_b(s_a, s_b) = \pi_b(r_a, s_b)$. And likewise with a and b interchanged.

If a game satisfies NRT, it also satisfies TDI. Yet many games of interest satisfy TDI, but fail to satisfy NRT. For example, zero sum games satisfy TDI, but may fail to satisfy NRT.

Proposition 3. *(i) Fix a PI game* Γ *that satisfies TDI. If* $Q_a \times Q_b$ *is an EFBRS, then there exists a pure-strategy Nash equilibrium, viz.* (s_a, s_b)*, so that each profile in* $Q_a \times Q_b$ *is outcome equivalent to* (s_a, s_b)*.*

(ii) Fix a PI game Γ *that satisfies NRT. If* (s_a, s_b) *is a pure-strategy Nash equilibrium in sequentially justifiable strategies, then there is an EFBRS, viz.* $Q_a \times Q_b$*, so that* $(s_a, s_b) \in Q_a \times Q_b$*.*

The proof can be found in Appendix D. Taken together, Theorem 1 and Proposition 3 give the following corollary.

Corollary 2. *(i) Fix a PI game* Γ *that satisfies TDI and has an epistemic type structure. If there is RCSBR at the state* (s_a, t_a, s_b, t_b)*, then* (s_a, s_b) *is outcome equivalent to a pure-strategy Nash equilibrium.*

(ii) Fix a PI game Γ *that satisfies NRT and has a pure-strategy Nash equilibrium, viz.* (s_a, s_b)*, in sequentially justifiable strategies. Then there exist an epistemic structure and a state thereof, viz.* (r_a, t_a, r_b, t_b)*, at which there is RCSBR and* $(r_a, r_b) = (s_a, s_b)$*.*

Why the connection between EFBRS's and Nash equilibria? Recall that if each player is "rational" (i.e., maximizes subjective expected utility) and places probability 1 on the actual strategy choices by the other player, then the strategy choices constitute a Nash equilibrium. In a PI game that satisfies TDI, RCSBR imposes a form of correct beliefs about the actual outcomes that obtain. Let us recast this at the level of the solution concept: In a PI game that satisfies TDI, each strategy profile in a given EFBRS is outcome equivalent. (This is Lemma 8 in Appendix D.) So along the path of play, the associated CPS('s) must assign probability 1 to a particular outcome—the outcome associated with the EFBRS, i.e., the "correct" outcome. (This uses condition (ii) of an EFBRS.) With this, we get a Nash outcome (but not necessarily the Nash strategies).[9]

This was the intuition for part (i) of Corollary 2. The proof closely follows the proof of Proposition 6.1a in Brandenburger and Friedenberg (2010), although now making use of the EFBRS properties. (The proof in Brandenburger and Friedenberg 2010 makes use of properties of self-admissible sets.)

[9] Ben-Porath (1997) gives another epistemic analysis of perfect-information games. His analysis is based on "rationality and common initial belief of rationality" plus a grain of truth assumption. It also gives Nash outcomes.

The converse, i.e., part (ii), is novel. (In particular, both the "no ties" condition and the proof are quite different from the analysis in Brandenburger and Friedenberg 2010.) A Nash equilibrium in sequentially justifiable strategies, in general, satisfies conditions (i) and (ii) of an EFBRS. However, it may fail the maximality criterion. Indeed, the proof makes use of all three properties of Definition 9; see Appendix D.

There is a gap between parts (i) and (ii) of Proposition 3. In particular, part (i) says that starting from an EFBRS, we can get a pure Nash outcome, while part (ii) says that starting from a sequentially justifiable pure-strategy Nash equilibrium, we can get an EFBRS.

We cannot improve part (ii) to say that starting from any pure Nash equilibrium, we get an EFBRS. (This is because a Nash equilibrium may not be sequentially justifiable; see Appendix D.) We do not know if we can improve part (i) to say that starting from an EFBRS, we get a pure-strategy Nash equilibrium in sequentially justifiable strategies. (Appendix D elaborates on this issue.) However, starting from an EFBRS, we can get a mixed-strategy Nash equilibrium that satisfies a "sequential justifiability" condition.

Consider a pure-strategy profile (s_a, s_b) and a mixed-strategy profile $(\varpi_a, \varpi_b) \in \mathcal{P}(S_a) \times \mathcal{P}(S_b)$. Call (s_a, s_b) and (ϖ_a, ϖ_b) *outcome equivalent* if $\pi(s_a, s_b) = \pi(\varpi_a, \varpi_b)$. Likewise, call $Q_a \times Q_b \subseteq S_a \times S_b$ and $(\varpi_a, \varpi_b) \in \mathcal{P}(S_a) \times \mathcal{P}(S_b)$ *outcome equivalent* if each $(s_a, s_b) \in Q_a \times Q_b$ is outcome equivalent to (ϖ_a, ϖ_b).

PROPOSITION 4. *Fix a PI game that satisfies TDI. If $Q_a \times Q_b$ is an EFBRS, then there exists a mixed-strategy Nash equilibrium, viz.* (σ_a, σ_b), *so that*

(i) $Q_a \times Q_b$ is outcome equivalent to (σ_a, σ_b) and

(ii) each $s_a \in \operatorname{Supp} \sigma_a$ (resp. $s_b \in \operatorname{Supp} \sigma_b$) is sequentially justifiable.

Proposition 4 gives that if we begin with an EFBRS, we can construct an equivalent mixed-strategy Nash equilibrium. The Nash equilibrium has the property that each strategy in its support is sequentially justifiable. But it is important to note that this does not necessarily give that the mixed-strategy itself is sequentially justifiable.[10] More to the point, given a PI game that satisfies TDI and some mixed-strategy Nash equilibrium, viz. (σ_a, σ_b), does there exist some pure-strategy Nash equilibrium, viz. (s_a, s_b), so that s_a (resp. s_b) is contained in the support of σ_a (resp. σ_b)? If so, using Proposition 4, we get that starting from an EFBRS, there is a pure-strategy Nash equilibrium in sequentially justifiable strategies. But this too is not known.

9. DISCUSSION

In this section, we discuss some conceptual aspects of the paper as well as some extensions.

[10]In non-PI games, we can construct a mixed-strategy Nash equilibrium, viz. (σ_a, σ_b), where each strategy in the support of σ_a and σ_b is sequentially justifiable, but σ_a is itself not sequentially justifiable. The question remains whether the same can occur in PI games.

9.a *Context-dependent forward induction*

We characterize the behavioral implications of forward induction reasoning across all type structures. Why the interest in such a result?

When we analyze a strategic situation, we specify the game (matrix or tree). But, in practice, there is a context to the strategic situation studied—e.g., players come to the game with social conventions, a history, etc. This context influences what beliefs players do vs. do not consider possible. If this is the case, it may be of interest to study a given game relative to different type structures, depending on the context within which the game is played.

One case of particular interest is where the analyst does not know the context, i.e., does not know which beliefs are vs. are not "transparent" to the players. If this is the case, the analyst will want to understand the behavioral implications of forward induction reasoning across all type structures. By Theorem 1, he should apply the EFBRS concept. (Contrast this with extensive-form rationalizability: The analyst should apply the extensive-form rationalizability concept, if he is interested in forward induction reasoning and understands that the players consider all possible beliefs. This is the implication of Proposition 6 in Battigalli and Siniscalchi 2002.)

9.b *Restrictions on beliefs*

In Section 9.a, we implicitly equated analyzing forward induction reasoning across all "transparent restrictions on players beliefs" with analyzing forward induction reasoning across all type structures. We can make this step precise. First, formalize the idea that certain (events about) beliefs are transparent to the players. For this, begin with Battigalli and Siniscalchi's (1999a) canonical construction of a type structure; this type structure contains all hierarchies of conditional beliefs (satisfying coherency and common belief of coherency). Let us look at the self-evident events within this structure. Loosely, we look at events $S_a \times E_a \times S_b \times E_b \in \mathcal{B}(S_a \times T_a \times S_b \times T_b)$ such that whenever $E = S_a \times E_a \times S_b \times E_b$ obtains, there is "common belief of E" in the following sense: each player assigns probability 1 to E at each node, each player assigns probability 1 at each node to the other player assigning probability 1 to E at each node, etc.[11] These self-evident events represent transparent restrictions on players' beliefs. Each type structure can be mapped into the canonical construction and, in a certain sense, each type structure forms a self-evident event in the canonical construction, i.e., under this mapping.[12] Furthermore, each such self-evident event in the canonical type structure corresponds to a "smaller" type structure. Forward induction reasoning is preserved under these mappings. (There is an equivalence between rationality in the small structure and "rationality and the self-evident event" in the large structure, and similarly for strong belief; see Battigalli and Friedenberg 2009 for the formal statement.)

[11] This is equivalent to the requirement that at each state where $E = S_a \times E_a \times S_b \times E_b$ obtains, each player assigns probability 1 to E at each of his information sets.

[12] This statement presumes that the image of the type set (under the mapping to the canonical construction) is measurable.

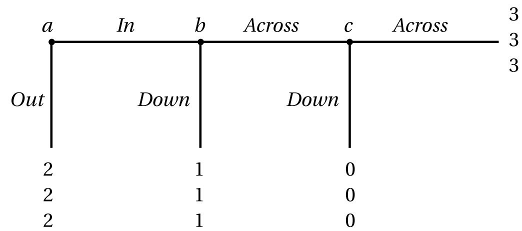

Figure 7. A modification of Figure 6.

There is a special type of transparent restriction on beliefs: those generated only by restrictions on first-order beliefs. In this case, there are explicit restrictions on first-order beliefs and the only restrictions on higher-order beliefs are those generated implicitly by the restrictions on first-order beliefs. (For instance, in the lady's choice convention, we explicitly restrict Bob's first-order beliefs, requiring that he assign probability 1 to Ann playing *Up*. This implicitly imposes a strong restriction on Ann's second-order beliefs, requiring that she assign probability 1 to the event "Bob assigns probability 1 to Ann playing *Up*" and so on; see Example 1.) The restrictions on first-order beliefs, viz. Δ, generate a particular type of self-evident event. Analyzing RCSBR within the associated type structure leads to the Δ-rationalizable strategy set. Indeed, this is related to Battigalli and Siniscalchi's (2003) motivation in defining directed rationalizability.[13]

9.c *Two versus three player games*

Here we have focused on two-player games. The main results (Theorem 1 and Corollary 1) extend to games with three or more players, up to issues of correlation. Specifically, if we allow for correlated assessments in Definition 8, then we must also allow for correlated assessments in Definition 9. A similar statement holds for the case of independence—although, of course, care is needed in defining independence for CPS's. The central issue is that Charlie's belief about Bob should not change after Charlie learns information only about Ann. (The idea dates back to Hammond 1987 and is related to the "do not signal what you do not know" condition of Fudenberg and Tirole 1991. See Battigalli 1996 for a formalization of the idea and a discussion of Fudenberg and Tirole 1991.)

There is an additional issue that arises in the three-player case: Should we require that Ann strongly believes "Bob and Charlie are rational" or should we instead require that Ann strongly believes "Bob is rational" and strongly believes "Charlie is rational"? Arguably, in the case of independence, we should require the latter.

How does this affect our analysis of games? Amend Figure 6, so that it is a three-player game as in Figure 7. Consider a state at which there is RCSBR in the sense explained above (i.e., Bob has an independent assessment and strongly believes both "Ann

[13]The treatment here is due to Battigalli and Prestipino (2011). It is related to, but somewhat different from, the epistemic assumptions of Battigalli and Siniscalchi (2003, 2007). It is important to note that under either treatment, an amendment is needed to Battigalli and Siniscalchi's (2003) definition of Δ-rationalizability; see footnote 6.

is rational" and "Charlie is rational"). Let us ask which strategies can be played. Of course, using rationality, Charlie must play *Across* (at this state). Next we require that a type of Bob strongly believes "Ann is rational" and also "Charlie is rational." So, conditional on Bob's information set being reached, this type must maintain a hypothesis Charlie is rational, and so that Charlie plays *Across*. In this case, Bob's unique best response is to play *In*. Turning to Ann, we see that under an RCSBR analysis, she chooses *In*. So we only get the backward induction outcome. (Battigalli and Siniscalchi 1999b provide a "context free" epistemic analysis of this notion of independent rationalization.)

This example also shows that in the case of independence, Proposition 3(ii) does not hold. If we instead consider the case of correlation, then it may also be natural to instead require that Bob strongly believes "Ann and Charlie are rational" (i.e., as opposed to strong belief of "Ann is rational" and strong belief of "Charlie is rational"). Of course, it may be the case that when Bob's node is reached, he must forgo the hypothesis that "Ann and Charlie are rational." Thus, in this case, we do have an analogue of Proposition 3(ii). Indeed, both parts (i) and (ii) of Proposition 3 hold for the case of correlation.

Appendix A: Proofs for Section 4

Proof of Property 1. Fix an event $F \in \mathcal{E}$ with $F \cap \bigcap_m E_m \neq \varnothing$. Then $F \cap E_m \neq \varnothing$ for all m. So for each m, $\mu(E_m|F) = 1$. (This is because μ strongly believes each E_m.) But then $\mu(\bigcap_m E_m|F) = 1$. □

Proof of Property 2. Fix an event $F \in \mathcal{E}$ with $F \cap \text{proj}_{\Omega_1} E \neq \varnothing$. Then $(F \times \Omega_2) \cap E \neq \varnothing$. Since, by assumption, $\text{proj}_{\Omega_1} E$ is Borel, $\text{marg}_{\Omega_1} \mu(\text{proj}_{\Omega_1} E|F)$ is well defined. Since μ strongly believes E, $\mu(E|F \times \Omega_2) = 1$. Then $(\text{marg}_{\Omega_1} \mu)(\text{proj}_{\Omega_1} E|F) = 1$, as required. □

Appendix B: Directed rationalizability

In the text, we argue that for each epistemic type structure, there is a set of first-order beliefs Δ so that the projection of the RCSBR set is the Δ-rationalizable strategy set. The purpose of this appendix is to illustrate that this set of first-order beliefs may not correspond to the set of all first-order beliefs allowed by the epistemic type structure.

Figure 8 is a game of battle of the sexes preceded by an observed "money burning" move by Bob. (See Ben-Porath and Dekel 1992.) Here, Ann and Bob are playing a BoS game. However, prior to the game, Bob has the option to *Burn* (*B*) or *Not Burn* (*NB*) $2.

Suppose society has formed a modified version of the lady's choice convention. Now, there are no restrictions on players' first-order beliefs. (So, in particular, there are types of Bob who think Ann does not go for her best payoff.) But there is a restriction on Ann's second-order beliefs. Specifically, conditional on observing so-called normal behavior (i.e., a decision to *Not Burn*), Ann thinks that Bob thinks she goes for her best payoff and chooses *Up*. There is no restriction on Ann's second-order belief conditional on

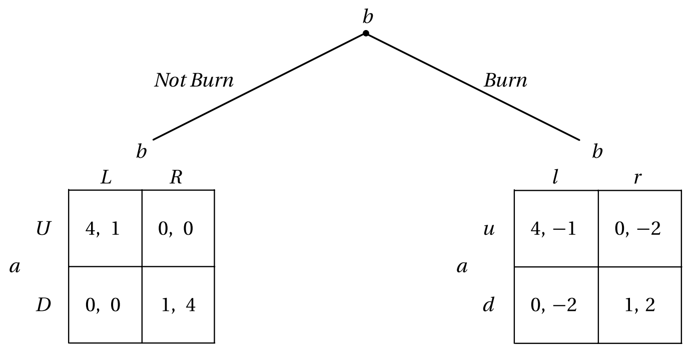

Figure 8. Battle of the sexes with money burning.

observing "strange" behavior, i.e., on observing a decision to *Burn*. Likewise, there are no restrictions on Bob's second-order beliefs, etc.

We can model this modified version of the lady's choice convention by a type structure $\langle S_a, S_b; \mathcal{S}_a, \mathcal{S}_b; T_a, T_b; \beta_a, \beta_b\rangle$ based on the game in Figure 8. Now, β_b is onto but β_a is not. Formally, write $[Up]_a$ for the event "Ann plays *Up*, if Bob does *Not Burn*," i.e., $[Up]_a = \{\textit{Up-down}, \textit{Up-up}\} \times T_a$, and write $[NB]_b$ for the event "Bob does *Not Burn*," i.e., $[NB]_b = \{\textit{NB-Left}, \textit{NB-Right}\} \times T_b$. Let U_b be the set of types $t_b \in T_b$ with $\beta_b(t_b)([Up]_a|S_a \times T_a) = 1$, i.e., the set of types of Bob that assign probability 1 to the event "Ann plays *Up*, when Bob chooses *Not Burn*." Then, for each type $t_a \in T_a$,

$$\beta_a(t_a)(S_b \times U_b|[NB]_b) = 1,$$

i.e., conditional on Bob choosing *Not Burn*, each type of Ann assigns probability 1 to the event that "Bob believes that 'Ann plays *Up*, when Bob does *Not Burn*.'" For any belief μ_a of Ann with $\mu_a(S_b \times U_b|[NB]_b) = 1$, there is a type t_a so that $\beta_a(t_a) = \mu_a$. (See Appendix A in Battigalli and Friedenberg 2009 on how to construct such a type structure.)

The set of first-order beliefs induced by this type structure is $\Delta = \mathcal{C}(S_b) \times \mathcal{C}(S_a)$. The Δ-rationalizable set is $\{\textit{Down-down}\} \times \{\textit{NB-Right}\}$. (This is also the set of extensive-form rationalizable strategies.) It is obtained as follows: In round one, the strategy *B-left* is dominated by *NB-Left*, but all other strategies (of both players) are optimal under some CPS. It follows that

$$S_a^{\Delta,1} \times S_b^{\Delta,1} = S_a \times \{\textit{NB-Left}, \textit{NB-Right}, \textit{B-right}\}.$$

But now note that the choice of *up* by Ann cannot be optimal under any CPS that strongly believes $\{\textit{NB-Left}, \textit{NB-Right}, \textit{B-right}\}$. (If a CPS strongly believes $\{\textit{NB-Left}, \textit{NB-Right}, \textit{B-right}\}$, then conditional on *Burn* being played, the CPS must assign probability 1 to *right*, in which case *up* is not a best response.) So

$$S_a^{\Delta,2} \times S_b^{\Delta,2} = \{\textit{Up-down}, \textit{Down-down}\} \times S_b^{\Delta,1}.$$

Turning to Bob, if a CPS strongly believes $\{\textit{Up-down}, \textit{Down-down}\}$, then *B-right* yields an expected payoff of 2 and *NB-Left* yields an expected payoff of at most 1. So

$$S_a^{\Delta,3} \times S_b^{\Delta,3} = S_a^{\Delta,2} \times \{\textit{NB-Right}, \textit{B-right}\}.$$

Now, if a CPS strongly believes $\{NB\text{-}Right, B\text{-}right\}$, *Down-down* is the only sequentially optimal strategy, so

$$S_a^{\Delta,4} \times S_b^{\Delta,4} = \{Down\text{-}down\} \times S_b^{\Delta,3}.$$

Finally, if a CPS strongly believes $\{Down\text{-}down\}$, *NB-Right* is the only sequentially optimal strategy, so

$$S_a^{\Delta,5} \times S_b^{\Delta,5} = \{Down\text{-}down\} \times \{NB\text{-}Right\}.$$

But the projection of event RCSBR onto $S_a \times S_b$ is $\{Up\text{-}down\} \times \{B\text{-}right\}$. It is obtained as follows. In round one, for each belief about the strategies of the other player, there is a type that holds that belief. So, here too,

$$\text{proj}_{S_a} R_a^1 \times \text{proj}_{S_b} R_b^1 = S_a \times \{NB\text{-}Left, NB\text{-}Right, B\text{-}right\}.$$

Now consider a type t_a that strongly believes R_b^1. Recall that, conditional on Bob choosing not to burn, each type of Ann assigns probability 1 to the event that "Bob believes that 'Ann plays *Up*, when Bob does not burn.'" So if t_a strongly believes R_b^1, it must assign zero probability to $\{NB\text{-}Right\} \times T_b$. For such a type t_a, $(Down\text{-}down, t_a)$ is irrational. So

$$\text{proj}_{S_a} R_a^2 \times \text{proj}_{S_b} R_b^2 = \{Up\text{-}down\} \times \text{proj}_{S_b} R_b^1.$$

But now, if (s_b, t_b) is rational and t_b strongly believes R_a^2, then $s_b = B\text{-}right$, and so

$$\text{proj}_{S_a} R_a^3 \times \text{proj}_{S_b} R_b^3 = \{Up\text{-}down\} \times \{B\text{-}right\}.$$

Why the difference between the two approaches? We began with an epistemic structure and used the structure itself to form the set of first-order beliefs $\Delta = \mathcal{C}(S_b) \times \mathcal{C}(S_a)$. (So for each $\mu_a \in \Delta_a = \mathcal{C}(S_b)$, there is type $t_a \in T_a$ such that the marginal of $\beta_a(t_a)$ on S_b is μ_a, and likewise for b.) With this set of first-order beliefs, the strategies that survive one round of Δ-rationalizability are exactly the strategies that are consistent with rationality. But in the next round, we lose the equivalence: If $\beta_a(t_a)$ strongly believes R_b^1, then the marginal of $\beta_a(t_a)$ must strongly believe $S_b^{\Delta,1} = \text{proj}_{S_b} R_b^1$. (Here, we use the marginalization property of strong belief.) Thus $\text{proj}_{S_a} R_a^2 \subseteq S_a^{\Delta,2}$. But the converse does not hold. We have $Down\text{-}down \in S_a^{\Delta,2}$, but $Down\text{-}down \notin \text{proj}_{S_a} R_a^2$. The reason is that, conditional on Bob choosing *NB*, each $\beta_a(t_a)$ assigns probability 1 to the event "Bob assigns probability 1 to $[Up]_a$." So if Bob does not burn, Ann can only maintain a hypothesis that Bob is rational if she assigns probability 1 to Bob's playing *NB-Left*, in which case the choice *Down* is not a best response. With this, $S_a^{\Delta,2} = \{Up\text{-}down, Down\text{-}down\}$ and $\text{proj}_{S_a} R_a^2 = \{Up\text{-}down\}$. As a result, $S_b^{\Delta,3} = \{NB\text{-}Right, B\text{-}right\}$ and $\text{proj}_{S_b} R_b^3 = \{B\text{-}right\}$. It follows that $S_a^{\Delta,4} = \{Down\text{-}down\}$, despite the fact that $\text{proj}_{S_a} R_a^4 = \{Up\text{-}down\}$. The key to this last step is that *Up-down* is optimal under a CPS that strongly believes $\text{proj}_{S_b} R_b^3 \subsetneq S_b^{\Delta,3}$, but not optimal under a CPS that strongly believes $S_b^{\Delta,3}$. This can occur because strong belief fails a monotonicity requirement.

Appendix C: Examples and proofs for Section 7

We begin by showing that for the finitely repeated prisoner's dilemma, any EFBRS results in the *Defect-Defect* path of play. To show this, we need to make use of certain properties of EFBRS's. We again make use of these properties in Appendix D. We begin with the best response property.

Definition 15. Say $Q_a \times Q_b \subseteq S_a \times S_b$ satisfies the *best response property* if, for each $s_a \in Q_a$, there is a CPS $\mu_a \in \mathcal{C}(S_b)$, so that $s_a \in \rho_a(\mu_a)$ and μ_a strongly believes Q_b, and similarly for b.

An EFBRS satisfies the best response property. But the converse need not hold, i.e., $Q_a \times Q_b$ may satisfy the best response property, but fail to be an EFBRS because it violates the maximality condition. (See the example in Section 5.)

Let us introduce some notation to relate the whole game to its parts. Fix a game Γ and a subgame Σ. Write H_a^Σ for the set of a's information sets that are contained in Σ. We abuse notation and write $S_a(\Sigma)$ for the set of strategies of Γ that allow Σ. We also write $S_a^\Sigma = \prod_{h \in H_a^\Sigma} C_a(h)$ for the set of strategies of a in the subgame Σ. Each strategy $s_a^\Sigma \in S_a^\Sigma$ can be viewed as the projection of a strategy $s_a \in S_a(\Sigma)$ into S_a^Σ. Given a set $E_a \subseteq S_a$, write E_a^Σ for the set of strategies $s_a^\Sigma \in S_a^\Sigma$ so that there is some $s_a \in E_a \cap S_a(\Sigma)$ whose projection into S_a^Σ is s_a^Σ. We write π_a^Σ and π_b^Σ for the payoff functions associated with the subtree Σ. So if (s_a, s_b) allows Σ, then $\pi^\Sigma(s_a^\Sigma, s_b^\Sigma) = \pi(s_a, s_b)$.

Lemma 1. *Fix a game Γ and a subgame Σ. If $Q_a \times Q_b$ satisfies the best response property for the game Γ, then $Q_a^\Sigma \times Q_b^\Sigma$ satisfies the best response property for the subgame Σ.*

Proof. If $Q_a^\Sigma \times Q_b^\Sigma = \varnothing$ (if no profile in $Q_a \times Q_b$ allows Σ), then it is immediate that $Q_a^\Sigma \times Q_b^\Sigma$ satisfies the best response property. So we suppose $Q_a^\Sigma \times Q_b^\Sigma \neq \varnothing$.

Fix a strategy $s_a^\Sigma \in Q_a^\Sigma$. Then there exists a strategy $s_a \in Q_a \cap S_a(\Sigma)$ whose projection into $\prod_{h \in H_a^\Sigma} C_a(h)$ is s_a^Σ. Since $s_a \in Q_a$, we can find a CPS $\mu_a \in \mathcal{C}(S_b)$ so that $s_a \in \rho_a(\mu_a)$ and μ_a strongly believes Q_b.

Let $\mathcal{S}_b^\Sigma$ be the set of all $S_b^\Sigma(h)$ for $h \in H_a^\Sigma$. Given an event $E_b^\Sigma \subseteq S_b^\Sigma$, write $E_b \subseteq S_b$ for the set of all $s_b \in S_b(\Sigma)$ so that the projection of s_b into S_b^Σ is in E_b^Σ. Then, define $\nu_a^\Sigma(\cdot|\cdot) : \mathcal{B}(S_b^\Sigma) \times \mathcal{S}_b^\Sigma \to [0, 1]$ so that, for each event $E_b^\Sigma \subseteq S_b^\Sigma$ and each $S_b^\Sigma(h) \in \mathcal{S}_b^\Sigma$, $\nu_a^\Sigma(E_b^\Sigma | S_b^\Sigma(h)) = \mu_a(E_b | S_b(h))$. It is readily verified that ν_a^Σ is a CPS on $(S_b^\Sigma, \mathcal{S}_b^\Sigma)$.

Since s_a allows Σ and s_a is sequentially optimal under μ_a, it follows that s_a^Σ is sequentially optimal under ν_a^Σ. Fix some $S_b^\Sigma(h) \in \mathcal{S}_b^\Sigma$. If $Q_b^\Sigma \cap S_b^\Sigma(h) \neq \varnothing$, then $Q_b \cap S_b(h) \neq \varnothing$. So, in this case, $\nu_a^\Sigma(Q_b^\Sigma | S_b^\Sigma(h)) \geq \mu_a(Q_b | S_b(h)) = 1$. This establishes that ν_a^Σ strongly believes Q_b^Σ.

Interchanging a and b establishes the result. □

We use Lemma 1 to show the next lemma.

Lemma 2. *Consider the N-repeated prisoner's dilemma as given in Figure 4. If $Q_a \times Q_b$ satisfies the best response property for this game, then each strategy profile in $Q_a \times Q_b$ results in the Defect-Defect path.*

Proof. The proof very closely follows the proof of Example 3.2 in Brandenburger and Friedenberg (2010). It is by induction on N. For $N = 1$, the result is immediate. Assume the result holds for some N and we show it holds for $N + 1$.

Consider some $Q_a \times Q_b$ of the $N + 1$ repeated prisoner's dilemma that satisfies the best response property. Suppose there is a strategy $s_a \in Q_a$ that plays *Cooperate* in the first period. Fix a strategy $s_b \in Q_b$. If s_b plays *Cooperate* (resp. *Defect*) in the first period, Ann gets c (resp. e) in the first period. By Lemma 1 and the induction hypothesis, Ann gets a payoff of zero in periods $2, \ldots, N$. So for each s_b in Q_b, $\pi_a(s_a, s_b) = c$ if s_b plays *Cooperate* in the first period and $\pi_a(s_a, s_b) = e$ if s_b plays *Defect* in the first period.

Now, instead, consider the strategy r_a that plays *Defect* in every period, irrespective of the history. Again, fix a strategy $s_b \in Q_b$. If s_b plays *Cooperate* in the first period, then $\pi_a(r_a, s_b) \geq d$, and if $s_b \in Q_b$ plays *Defect* in the first period, then $\pi_a(r_a, s_b) \geq 0$.

Putting the above together gives that under any CPS that strongly believes Q_b, we must have that r_a is a strictly better response than $s_a \in Q_a$ at the first information set. But this contradicts $Q_a \times Q_b$ satisfying the best response property. □

Corollary 3. *Consider the N-repeated prisoner's dilemma as given in Figure 4. If $Q_a \times Q_b$ is an EFBRS, then each strategy profile in $Q_a \times Q_b$ results in the Defect-Defect path.*

Now we turn to Proposition 2. We show the result for a somewhat more general set of games, i.e., games where, in a sense, the information structure is determined by the subgames.

Definition 16. Fix a game Γ. Say a subgame Σ is *sufficient* for an information set $h \in H$ if h is contained in Σ and the set of strategy profiles that allow Σ is exactly $S_a(h) \times S_b(h)$.

Notice that there may be two subgames, viz. Σ and $\bar{\Sigma}$, that are sufficient for h.[14] If so, either Σ is a subgame of $\bar{\Sigma}$ or $\bar{\Sigma}$ is a subgame of Σ. When there are two subgames that are sufficient for h, we typically are interested in the *last subgame* Σ *sufficient for* h, i.e., so that no proper subgame of Σ is sufficient for h. Also notice that there may be no subgame that is sufficient for an information set h. Refer to the game in Figure 5. There the only subgame is the entire game. But this subgame is not sufficient for the information set, viz. h, at which Bob moves. To see this, notice that the strategy $s_a = Out$ (trivially) allows the subgame, but does not allow h.

Definition 17. Say a game Γ is *determined by its subgames* if, for each information set $h \in H$, there is a subgame Σ that is sufficient for h.

[14] This may happen if there is a node x where no player is active, i.e., $C_a(x)$ and $C_b(x)$ are singletons.

The game in Figure 5 is not determined by its subgame; as we have seen, there is no subgame that is sufficient for the information set at which Bob moves. Battigalli and Friedenberg (2009) characterize Definition 17 in terms of primitives of the game (as opposed to a condition about strategies).

Before stating the generalization of Proposition 2, we need to extend the definition of NRT to cover games with imperfectly observable actions.

Definition 18. Fix two distinct terminal nodes $z = (x, c^1, \dots, c^K)$ and $z' = (x, d^1, \dots, d^L)$. Say *a is decisive for* (z, z') if the following conditions hold.

(i) $c_a^1 \neq d_a^1$,

(ii) $c_b^1 = d_b^1$, and

(iii) if $(x, c^1, \dots, c^k)$ and $(x, d^1, \dots, d^l)$ are in the same information set for b, then $c_b^{k+1} = d_b^{l+1}$.

The idea is that a is decisive for $(z, z') = ((x, c^1, \dots, c^K), (x, d^1, \dots, d^L))$ if a is the only player who determines which of the two terminal histories occurs. So a moves at the last common predecessor of z and z', viz. x, and makes distinct choices at this node, i.e., $c_a^1 \neq d_a^1$. But b's choice along these paths does not determine which of z vs. z' occurs. So b makes the same choice whenever he cannot observe a's choice among c_a^1 vs. d_a^1.

Remark 3. If the game has observable actions, then a is decisive for $(z, z') = ((x, c^1, \dots, c^K), (x, d^1, \dots, d^L))$ if and only if $c_a^1 \neq d_a^1$ and $c_b^1 = d_b^1$.

Definition 19 (Battigalli 1997). A game satisfies *no relevant ties* (NRT) if whenever a (resp. b) is decisive for (z, z'), $\Pi_a(z) \neq \Pi_a(z')$.

Now, here is the generalization of Proposition 2.

Proposition 5. *Fix a game* Γ *that is determined by its subgames and a pure-strategy SPE, viz.* (s_a, s_b).

(i) There is an EFBRS, viz. $Q_a \times Q_b$*, so that* $[s_a] \times [s_b] \subseteq Q_a \times Q_b$.

(ii) If Γ *satisfies NRT, then* $[s_a] \times [s_b]$ *is an EFBRS.*

Before coming to the proof, it is useful to record some facts about games determined by their subgames. Fix a pure-strategy SPE, viz. (s_a, s_b), of a game Γ determined by its subgames. Construct maps $f_a : H \to S_a$ and $f_b : H \to S_b$ that depend on this SPE. To do so, fix some $h \in H$ and let Σ be the last subgame sufficient for h. Write x for the root of subgame Σ (which may be Γ itself). If $\Sigma = \Gamma$, set $f_a(h) = s_a$. If Σ is a proper subtree of Γ, then we can write $x = (c^1, \dots, c^K)$. In this case, let $f_a(h)$ be the strategy that (i) chooses c_a^1 at $\{\phi\}$, (ii) chooses c_a^k at an information set that contains $(c^1, \dots, c^{k-1})$, i.e., an initial segment of $(c^1, \dots, c^K)$, and (iii) makes the same choice as s_a at all other information sets. So if s_a allows h, then $f_a(h) = s_a$. Also, $f_a(h)$ is well defined and allows h precisely because Γ is determined by its subgames. (Again, refer to the game in Figure 5, and take h to be the information set at which Bob moves. Consider the SPE $(s_a, s_b) = (Out, Right)$. Then $f_a(h) = Out$, which precludes h.)

Write $S(h)$ for the set of strategy profiles that allow an information set h. In games determined by their subgames, there is a natural order on sets of the form $S(h)$ for $h \in H$. Specifically, for any pair of information sets h and i (in H), either $S(h) \subseteq S(i)$, $S(i) \subseteq S(h)$, or $S(h) \cap S(i) = \varnothing$.[15] To see this, let Σ_h (resp. Σ_i) be sufficient for h (resp. i). We have that either Σ_h is a subgame of Σ_i, Σ_i is a subgame of Σ_h, or they are disjoint subgames. With this, the order follows from the definition of sufficiency. If $S(h) \subseteq S(i)$, say h *follows* i. Say h and i are *ordered* if either h follows i or i follows h. Say h and i are *unordered* otherwise, i.e., if $S(h) \cap S(i) = \varnothing$.

The proofs of the following results are immediate.

Lemma 3. *Fix a game Γ that is determined by its subgames. Also fix some SPE, viz. (s_a, s_b). Construct (f_a, f_b) as above. If $f_a(h)$ allows i, and either h and i are unordered or i follows h, then $f_a(i) = f_a(h)$.*

Lemma 4. *Fix a game Γ that is determined by its subgames and some SPE (s_a, s_b). For each $h \in H_a$,*

$$\pi_a(f_a(h), f_b(h)) \geq \pi_a(r_a, f_b(h)) \quad \text{for all } r_a \in S_a(h).$$

Lemma 5. *Fix some $\mu_a \in \mathcal{C}(S_b)$. If $s_a \in \rho_a(\mu_a)$, then $[s_a] \subseteq \rho_a(\mu_a)$.*

Proof of Proposition 5. Fix a pure-strategy SPE, viz. (s_a, s_b). Construct maps $f_a : H \to S_a$ and $f_b : H \to S_b$ as above. We use these maps to construct CPS's $\mu_a \in \mathcal{C}(S_b)$ and $\mu_b \in \mathcal{C}(S_a)$. Specifically, set $\mu_a(f_b(h)|S_b(h)) = 1$ for each $h \in H_a$. And likewise for a and b interchanged.

First we show that μ_a is indeed a CPS. It is immediate that μ_a satisfies conditions (i) and (ii) of Definition 1. For condition (iii), fix information sets $h, i \in H_a$ so that $S_b(i) \subseteq S_b(h)$. If $f_b(h) \in S_b(i)$, then $f_b(i) = f_b(h)$ (Lemma 3). So for each event $E \subseteq S_b(i)$,

$$\mu_a(E|S_b(h)) = \mu_a(E|S_b(i)) \times 1 = \mu_a(E|S_b(i))\mu_a(S_b(i)|S_b(h)).$$

If $f_b(h) \notin S_b(i)$, then for each event $E \subseteq S_b(i)$,

$$\mu_a(E|S_b(h)) = 0 = \mu_a(E|S_b(i)) \times 0 = \mu_a(E|S_b(i))\mu_a(S_b(h)|S_b(i)),$$

as required. And likewise for b.

Now let $Q_a = \rho_a(\mu_a)$, i.e., the set of all strategies r_a that are sequentially optimal under μ_a, and likewise set $Q_b = \rho_b(\mu_b)$. We show that $Q_a \times Q_b$ is an EFBRS.

Fix some $r_a \in Q_a$. We show that r_a and μ_a jointly satisfy conditions (i)–(iii) of an EFBRS. In fact, it is immediate that conditions (i) and (iii) are satisfied, so we show condition (ii), i.e., that μ_a strongly believes Q_b.

[15]Note that in all perfect recall games, whenever $h, i \in H_a$, either $S(h) \subseteq S(i)$, $S(i) \subseteq S(h)$, or $S(h) \cap S(i) = \varnothing$. Here we have an analogous statement, when $h \in H_a$ and $i \in H_b$.

Fix an information set $h \in H_a$ with $Q_b \cap S_b(h) \neq \varnothing$. We show that $f_b(h) \in Q_b$, so that $\mu_a(Q_b|S_b(h)) = 1$. To show that $f_b(h) \in Q_b$, it suffices to show that for each information set $i \in H_b$ allowed by $f_b(h)$,

$$\pi_b(f_a(i), f_b(h)) \geq \pi_b(f_a(i), r_b) \quad \text{for all } r_b \in S_b(i). \tag{C.1}$$

Note that if either i follows h or h and i are unordered, then $f_b(h) = f_b(i)$. In either case, we can apply Lemma 4 to the information set i and get the desired result. So we focus on the case where h follows i.

Take $S(h) \subseteq S(i)$. Since $Q_b \cap S_b(h) \neq \varnothing$, there is a strategy $r_b \in Q_b \cap S_b(h)$. For this strategy r_b, we have that $\pi_b(f_a(i), r_b) \geq \pi_b(f_a(i), f_b(h))$, because r_b is sequentially optimal under μ_b, $\mu_b(f_a(i)|S_a(i)) = 1$, and $f_b(h) \in S_b(h) \subseteq S_b(i)$. We show that $\pi_b(f_a(i), r_b) = \pi_b(f_a(i), f_b(h))$, establishing (C.1).

Suppose, contra hypothesis, that $\pi_b(f_a(i), r_b) > \pi_b(f_a(i), f_b(h))$. Consider the information set j, so that the last common predecessor of $\zeta(f_a(i), r_b)$ and $\zeta(f_a(i), f_b(h))$ is contained in j. Now use the fact that r_b and $f_b(h)$ both allow h to get that either j follows h or j and h are unordered. In these cases, we have that $\pi_b(f_a(j), f_b(h)) \geq \pi_b(f_a(j), r_b)$. (This was established in the previous paragraph.) But now notice that, since either j follows h or j and h are unordered, we also have that either j follows i or j and i are unordered. In either case, using the fact that $f_a(i)$ allows j, we have $f_a(i) = f_a(j)$ (Lemma 3). So putting the above facts together, we get

$$\begin{aligned} \pi_b(f_a(i), f_b(h)) &= \pi_b(f_a(j), f_b(h)) \\ &\geq \pi_b(f_a(j), r_b) \\ &= \pi_b(f_a(i), r_b) \geq \pi_b(f_a(i), f_b(h)). \end{aligned}$$

But this contradicts the assumption that $\pi_b(f_a(i), r_b) > \pi_b(f_a(i), f_b(h))$.

We have established that $Q_a \times Q_b = \rho_a(\mu_a) \times \rho_b(\mu_b)$ is an EFBRS. By construction, $(s_a, s_b) \in \rho_a(\mu_a) \times \rho_b(\mu_b)$. So using Lemma 5, $[s_a] \times [s_b] \subseteq Q_a \times Q_b$. Now suppose the game tree has NRT. We show that if $(r_a, r_b) \in Q_a \times Q_b$, then $(r_a, r_b) \in [s_a] \times [s_b]$.

Fix some strategy $r_a \notin [s_a]$. Then there exists some $r_b \in S_b$ with $\zeta(s_a, r_b) \neq \zeta(r_a, r_b)$. Consider the last common predecessor of $\zeta(s_a, r_b)$ and $\zeta(r_a, r_b)$, viz. x, and let h be the information set that contains this node. Then there exists $(c^1, \dots, c^K)$ and $(d^1, \dots, d^L)$ so that $\zeta(s_a, r_b) = (x, c^1, \dots, c^K)$, $\zeta(r_a, r_b) = (x, d^1, \dots, d^L)$. Clearly, $c_a^1 = s_a(h) \neq r_a(h) = d_a^1$ and $c_b^k = r_b(h') = d_b^l$ whenever $(x, c^1, \dots, c^{k-1}), (x, d^1, \dots, d^L) \in h' \in H_b$. So a is decisive for $(\zeta(s_a, r_b), \zeta(r_a, r_b))$.

Now, by the analysis above, we have that $\pi_a(s_a, f_b(h)) \geq \pi_a(r_a, f_b(h))$. NRT says that, in fact, $\pi_a(s_a, f_b(h)) > \pi_a(r_a, f_b(h))$. This implies that $r_a \notin Q_a$, as required. □

Lemma 6. *If* Γ *has observable actions, then* Γ *is determined by its subgames.*

Proof. Fix an information set h. Since Γ has observable actions, $h = \{x\}$ for some node/history x. Now consider a node y that follows x. Then by observable actions, y is contained in the information set $\{y\}$. It follows that there is a subgame whose initial

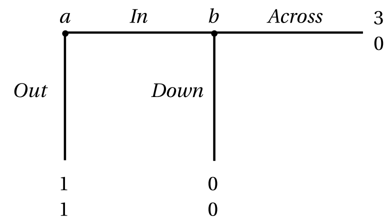

Figure 9. A PI game with relevant ties.

node is x, written Σ. Moreover, the set of strategies that allow Σ is exactly $S_a(h) \times S_b(h)$. So Γ is determined by its subgames. □

The proof of Proposition 2 is immediate from Proposition 5 and Lemma 6.

Finally, we conclude by pointing out the need for NRT in Proposition 5(ii).

Example 12. Figure 9 gives a game that fails NRT. Since it is a perfect-information game, it is determined by its subgames. Here, $(In, Across)$ is a pure-strategy SPE, but $\{In\} \times \{Across\}$ is not an EFBRS.

There is an EFBRS, viz. $Q_a \times Q_b$, with $\{In\} \times \{Across\} \subseteq Q_a \times Q_b$, e.g., $\{In\} \times \{Across, Down\}$. (Of course, part (i) of Proposition 2 says there must be some such EFBRS.) But every EFBRS, viz. $Q_a \times Q_b$, must have $Q_b = \{Across, Down\}$. (Here we use condition (iii) of an EFBRS.) So $\{In\} \times \{Across\}$ is not an EFBRS. ◊

Appendix D: Examples and proofs for Section 8

In this appendix, we prove Propositions 3 and 4. We also provide examples to better understand the results.

D.I *No ties and Proposition 3*

Part (i) of Proposition 3 requires TDI and part (ii) of Proposition 3 requires NRT. Example 13 explains why part (i) requires TDI.

Example 13. Return to Example 12, which fails TDI. There we see that $(In, Down)$ is contained in an EFBRS. But it is not outcome equivalent to a pure-strategy Nash equilibrium. ◊

Observe that when Bob moves, he is indifferent between *In* and *Out*. Now turn to a type of Ann that strongly believes Bob is rational. This type has a correct belief about what Bob's payoff will be if she plays *In*. But because the game fails TDI, she may have an incorrect belief about what her own payoff will be if she plays *In*. As such, a Nash outcome need not obtain.

Example 14 explains why we cannot replace NRT with the (weaker) TDI condition in part (ii) of Proposition 3.

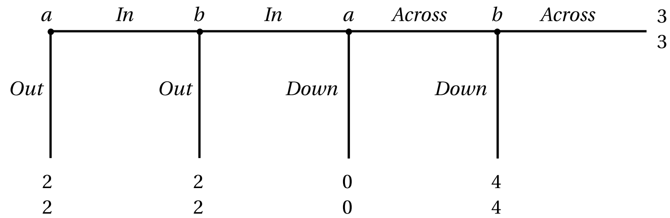

Figure 10. A game with TDI that fails NRT.

Example 14. Consider the game in Figure 10, which satisfies TDI, but violates NRT. Here, (Out, Out) is a Nash equilibrium in sequentially justifiable strategies. But if $Q_a \times Q_b$ is a (nonempty) EFBRS, then $Q_a \times Q_b = \{In\text{-}Across\} \times \{In\text{-}Down\}$. To see this, let $Q_a \times Q_b \neq \varnothing$ be an EFBRS. In this case, $Q_a \subseteq \{Out, In\text{-}Across\}$ and $Q_b \subseteq \{Out, In\text{-}Down\}$. (The strategy *In-Down* for Ann is dominated at her second information set, and the strategy *In-Across* for Bob is dominated at his second information set.) Also, *In-Across* is a weakly dominant strategy for Ann. So condition (iii) of an EFBRS implies that $In\text{-}Across \in Q_a$. It follows that if μ_b strongly believes Q_a, then μ_b must assign probability 1 to *In-Across* conditional on the event "Ann plays *In*." So *In-Down* is Bob's only strategy that is sequentially optimal given a CPS that strongly believes Q_a. This implies that $Q_b = \{In\text{-}Down\}$ and so $Q_a = \{In\text{-}Across\}$. ◊

In the above example, $\{(Out, Out)\}$ is disjoint from any EFBRS. While it satisfies conditions (i) and (ii) of an EFBRS, it fails condition (iii): If (Out, Out) is played, Ann gets a payoff of 2. But by going *In*, she can also assure herself an expected payoff of at least 2. As such, condition (iii) requires that we include *In-Across*.

To better understand what is going on, let us recast this at the epistemic level: If (Out, t_a) is rational, so is $(In\text{-}Across, t_a)$. With this, if Bob strongly believes that Ann is rational, then when his first information set is reached, he must maintain a hypothesis that Ann is playing *In-Across*; that is, he must maintain a hypothesis that Ann is playing a particular strategy that is not in $Q_a = \{Out\}$. As such, *Out* cannot be a best response for Bob.

The key is that the rationality of (Out, t_a) has implications for Ann's rationality at information sets precluded by *Out*. Notice that this happens because Ann is indifferent between the terminal nodes reached by (Out, Out) and $(In\text{-}Across, Out)$. (If Ann's payoffs from $(In\text{-}Across, Out)$ are strictly less than 2, (Out, t_a) can be rational without $(In\text{-}Across, t_a)$ being rational. Similarly, if Ann's payoffs from $(In\text{-}Across, Out)$ are strictly greater than 2, then (Out, Out) would not be a Nash equilibrium.) This is where the NRT condition comes in—it says that if Ann is decisive between two terminal nodes (as she is here), then she cannot be indifferent between those nodes.

D.II *Proof of Proposition 3(i)*

The proof follows immediately from the following lemma.

LEMMA 7. *Fix a perfect-information game that satisfies TDI. If $Q_a \times Q_b$ satisfies the best response property, then each $(s_a, s_b) \in Q_a \times Q_b$ is outcome equivalent to a Nash equilibrium.*

The proof of this lemma closely follows the proof of Proposition 6.1a in Brandenburger and Friedenberg (2010). It is by induction on the length of the tree. Specifically, fix a game Γ and a subgame Σ. The induction hypothesis states that if a set satisfies the best response property on Σ, then it is outcome equivalent to some Nash equilibrium. We know that if a set $Q_a \times Q_b$ satisfies the best response property on Γ, it also satisfies the best response property on the subgame Σ. (This is Lemma 1.) So if we fix a set that satisfies the best response property on the whole tree, then, by the induction hypothesis, it is outcome equivalent to a Nash equilibrium on each reached subgame. The proof uses this fact to construct a pure-strategy Nash equilibrium on the whole tree that is outcome equivalent to each profile in $Q_a \times Q_b$.

DEFINITION 20. Call $Q_a \times Q_a \subseteq S_a \times S_b$ a *constant set* if, for each $(s_a, s_b), (r_a, r_b) \in Q_a \times Q_b$, $\pi(s_a, s_b) = \pi(r_a, r_b)$.

LEMMA 8. *Fix a perfect-information game that satisfies TDI. If $Q_a \times Q_b$ satisfies the best response property, then $Q_a \times Q_b$ is a constant set.*

PROOF. The proof is by induction on the length of the tree. First, fix a tree of length 1 and suppose Ann moves at the initial node. Then Bob's strategy set is a singleton. So if $Q_a \times Q_b$ satisfies the best response property, then Ann is indifferent between each (s_a, s_b) and (r_a, s_b) in $Q_a \times Q_b$. By TDI, each profile in $Q_a \times Q_b$ is outcome equivalent.

Assume the result holds for any tree of length l or less. Fix a tree of length $l + 1$ and a set $Q_a \times Q_b$ satisfying the best response property. Suppose Ann moves at the initial node and can choose among nodes $n_1, \dots, n_K$. Each n_k can be identified with an information set and each is associated with a subgame $\Sigma = k$.

In particular, fix some subgame k with $Q_a^k \times Q_b^k \neq \varnothing$. Then $Q_a^k \times Q_b^k$ satisfies the best response property for the subgame k. (This is Lemma 1.) So by the induction hypothesis, $\pi^k(s_a^k, s_b^k) = \pi^k(r_a^k, r_b^k)$ for (s_a^k, s_b^k), $(r_a^k, r_b^k) \in Q_a^k \times Q_b^k$. Now note that for each $s_b \in Q_b$, $s_b^k \in Q_b^k$. (Here, we use the fact that Ann moves at the initial node.) Thus, given two strategies $s_a, r_a \in Q_a \cap S_a(\Sigma)$ and $s_b, r_b \in Q_b$, we have that $\pi(s_a, s_b) = \pi(r_a, r_b)$.

Now fix some $(s_a, s_b), (r_a, r_b) \in Q_a \times Q_b$, where $s_a \in S_a(k)$ and $r_a \in S_a(j)$. We have already established that $\pi(s_a, s_b) = \pi(r_a, r_b)$, for $k = j$. Suppose $k \neq j$. Since $s_a \in Q_a$, s_a is sequentially optimal under some $\mu_a(\cdot|\cdot)$ that strongly believes Q_b. So, in particular, s_a is optimal under $\mu_a(\cdot|S_b)$ with $\mu_a(Q_b|S_b) = 1$. With this,

$$\begin{aligned}\pi_a(s_a, s_b) &= \sum_{q_b \in Q_b} \pi_a(s_a, q_b)\mu_a(q_b|S_b) \\ &\geq \sum_{q_b \in Q_b} \pi_a(r_a, q_b)\mu_a(q_b|S_b) \\ &= \pi_a(r_a, r_b).\end{aligned}$$

(The first equality follows from the fact that for each $q_b \in Q_b$, $\pi_a(s_a, s_b) = \pi_a(s_a, q_b)$. This is a consequence of the last line in the preceding paragraph; likewise for the last equality.) By an analogous argument, $\pi_a(r_a, r_b) \geq \pi_a(s_a, s_b)$. So, $\pi_a(r_a, r_b) = \pi_a(s_a, s_b)$. By TDI, $\pi_b(r_a, r_b) = \pi_b(s_a, s_b)$. □

Proof of Lemma 7. The proof is by induction on the length of the tree. First, fix a tree of length 1 and suppose Ann moves at the initial node. Then Bob's strategy set is a singleton. The result follows from the fact that each $s_a \in Q_a$ is sequentially optimal under a CPS.

Now assume the result holds for any tree of length l or less. Suppose Ann moves at the initial node, and can choose among nodes $n^1, \ldots, n^K$. Each n^k can be identified with an information set and each is associated with a subgame $\Sigma = k$.

Fix some $(s_a, s_b) \in Q_a \times Q_b$ and suppose $s_a \in S_a(1)$. Note that $Q_a^1 \times Q_b^1$ satisfies the best response property (Lemma 1). So by the induction hypothesis, there is a Nash equilibrium of subgame 1, viz. (r_a^1, r_b^1), so that $\pi(s_a^1, s_b^1) = \pi(r_a^1, r_b^1)$. Consider a strategy $r_a \in S_a(1)$ so that the projection of r_a onto $\prod_{h \in H_a^1} C_a(h)$ is r_a^1. We need to show that we can choose $r_b^2, \ldots, r_b^K \in \bigtimes_{k=2}^K S_b^k$ so that, for each $q_a \in Q_a$ and associated $q_a^k \in S_a^k$, $\pi_a(r_a^1, r_b^1) \geq \pi_a(q_a^k, r_b^k)$. The profile $(r_a, (r_b^1, r_b^2, \ldots, r_b^K))$ is then a Nash equilibrium of the game.

Since $s_a \in Q_a$, there exists a CPS and an associated measure $\mu_a(\cdot|S_b)$ so that

$$\sum_{s_b \in S_b} [\pi_a(s_a, s_b) - \pi_a(q_a, s_b)]\mu_a(s_b|S_b) \geq 0$$

for all $q_a \in S_a$. Fix k from $2, \ldots, K$. Using Lemma 8,

$$\pi_a(r_a^1, r_b^1) = \pi_a(s_a^1, s_b^1) \geq \sum_{s_b^k \in S_k^b} \pi_a(q_a^k, s_b^k)(\operatorname{marg}_{S_b^k} \mu(\cdot|S_b))(s_b^k)$$

for any $q_a^k \in S_a^k$. Letting $(\overline{q}_a^k, \overline{q}_b^k) \in \arg\max_{S_a^k} \min_{S_b^k} \pi_a(\cdot, \cdot)$, we have in particular

$$\pi_a(r_a^1, r_b^1) \geq \sum_{s_b^k \in S_k^b} \pi_a(\overline{q}_a^k, s_b^k)(\operatorname{marg}_{S_k^b} \mu(\cdot|S_b))(s_b^k).$$

But $\pi_a(\overline{q}_a^k, q_b^k) \geq \pi_a(\overline{q}_a^k, \overline{q}_b^k)$ for any $q_b^k \in S_b^k$, by definition. So

$$\pi_a(r_a^1, r_b^1) \geq \sum_{s_b^k \in S_k^b} \pi_a(\overline{q}_a^k, \overline{q}_b^k)(\operatorname{marg}_{S_b^k} \mu(\cdot|S_b))(s_b^k) = \pi_a(\overline{q}_a^k, \overline{q}_b^k).$$

Set $(\underline{q}_a^k, \underline{q}_b^k) \in \arg\min_{S_b^k} \max_{S_a^k} \pi_a(\cdot, \cdot)$. By the minimax theorem for PI games (see, e.g., Ben-Porath 1997), $\pi_a(\overline{q}_a^k, \overline{q}_b^k) = \pi_a(\underline{q}_a^k, \underline{q}_b^k)$. It follows that $\pi_a(r_a^1, r_b^1) \geq \pi_a(\overline{q}_a^k, \overline{q}_b^k) = \pi_a(\underline{q}_a^k, \underline{q}_b^k)$. But $\pi_a(\underline{q}_a^k, \underline{q}_b^k) \geq \pi_a(q_a^k, \underline{q}_b^k)$ for any $q_a^k \in S_a^k$, by definition. So $\pi_a(r_a^1, r_b^1) \geq \pi_a(q_a^k, \underline{q}_b^k)$ for each $q_a^k \in S_a^k$. Setting each $r_b^k = \underline{q}_b^k$ gives the desired profile. □

D.III *Proof of Proposition 3(ii)*

Let us give the idea of the proof. We start with a set $Q_a \times Q_b = \{(s_a, s_b)\}$, where (s_a, s_b) is a pure Nash equilibrium in sequentially justifiable strategies. This set satisfies the best response property. (See Lemma 10 below.) In particular, the set Q_a is associated with a single CPS μ_a, satisfying the conditions of the best response property. We look at the set P_a of all strategies r_a that are sequentially optimal under μ_a. We use the fact that μ_a strongly believes Q_b (so assigns probability 1 to s_b at the initial information set) to get that Ann is indifferent between all outcomes associated with $P_a \times Q_b$. Indeed, by NRT, these strategy profiles must reach the same terminal node. Likewise, we define P_b and, using standard properties of a PI game tree, we get that all strategies in $P_a \times P_b$ reach the same terminal node.

So what have we done? We began with a set $Q_a \times Q_b$ and we expanded it to a set $P_a \times P_b$, with (i) $Q_a \times Q_b \subseteq P_a \times P_b$, (ii) all the profiles in $P_a \times P_b$ reach the same terminal node, and (iii) there is a CPS μ_a (resp. μ_b) that strongly believes Q_b (resp. Q_a) and such that P_a (resp. P_b) is the set of strategies that are sequentially optimal under $\mu_a(\cdot|\cdot)$ (resp. $\mu_b(\cdot|\cdot)$). We have successfully in constructed an EFBRS if the CPS μ_a (resp. μ_b) strongly believes P_b (resp. P_a) instead of Q_b (resp. Q_a). The key is that we can similarly expand $P_a \times P_b$ so that the new set satisfies similar properties. Since the game is finite, eventually the expanded set must coincide with the original set; that is, condition (i) must hold with equality. This gives the desired result.

Now we turn to the proof. First, we give a technical lemma.

Lemma 9. *Fix some $(\Omega, \mathcal{E})$ where Ω is finite. Let $\mu(\cdot|\cdot)$ be a CPS on $(\Omega, \mathcal{E})$ and let ϖ be a measure on Ω. Construct $\nu(\cdot|\cdot) : \mathcal{B}(\Omega) \times \mathcal{E} \to [0, 1]$ as follows: If $F \in \mathcal{E}$ with $\operatorname{Supp} \varpi \cap F \neq \emptyset$, then $\nu(\cdot|F) = \varpi(\cdot|F)$. Otherwise, $\nu(\cdot|F) = \mu(\cdot|F)$. Then $\nu(\cdot|\cdot)$ is a CPS.*

Proof. Let μ, ϖ, and ν be as in the statement of the lemma. Conditions (i) and (ii) of a CPS are immediate. Turn to condition (iii). For this, fix $E \in \mathcal{B}(\Omega)$ and $F, G \in \mathcal{E}$ with $E \subseteq F \subseteq G$.

First suppose that $\operatorname{Supp} \varpi \cap F \neq \emptyset$. Then

$$\begin{aligned} \nu(E|G) &= \frac{\varpi(E)}{\varpi(G)} \\ &= \frac{\varpi(E)}{\varpi(F)} \frac{\varpi(F)}{\varpi(G)} = \nu(E|F)\nu(F|G), \end{aligned}$$

where the first equality makes use of the fact that $E \subseteq G$, and the last equality makes use of the fact that $E \subseteq F$ and $F \subseteq G$. Next suppose that $\operatorname{Supp} \varpi \cap G = \emptyset$. Then $\operatorname{Supp} \varpi \cap F = \emptyset$, so that

$$\begin{aligned} \nu(E|G) &= \mu(E|G) \\ &= \mu(E|F)\mu(F|G) = \nu(E|F)\nu(F|G), \end{aligned}$$

as required. Finally, suppose that $\operatorname{Supp} \varpi \cap F = \emptyset$ but $\operatorname{Supp} \varpi \cap G \neq \emptyset$. Then

$$0 \leq \nu(E|G) \leq \nu(F|G) = \varpi(F|G) = 0,$$

where the last equality follows from the fact that Supp $\varpi \cap F = \emptyset$. Then

$$\nu(E|G) = 0$$
$$= \mu(E|F)\varpi(F|G) = \nu(E|F)\nu(F|G),$$

as required. □

LEMMA 10. *Let (s_a, s_b) be a Nash equilibrium in sequentially justifiable strategies. Then $\{(s_a, s_b)\}$ satisfies the best response property.*

PROOF. Let (s_a, s_b) be a Nash equilibrium in sequentially justifiable strategies. Then there exists a CPS $\mu_a(\cdot|\cdot)$ so that s_a is sequentially optimal under $\mu_a(\cdot|\cdot)$. Construct a CPS $\nu_b(\cdot|\cdot)$ so that $\nu_b(s_b|S_b(h)) = 1$ if $s_b \in S_b(h)$ and $\nu_b(\cdot|S_b(h)) = \mu_a(\cdot|S_b(h))$ otherwise. By Lemma 9, $\nu_b(\cdot|\cdot)$ is a CPS. It is immediate from the construction that s_a is sequentially optimal under $\nu_b(\cdot|\cdot)$ and that $\nu_b(\cdot|\cdot)$ strongly believes $\{s_b\}$, and, similarly with a and b reversed. □

DEFINITION 21. Fix a constant set $Q_a \times Q_a \subseteq S_a \times S_b$. Call $P_a \times P_a \subseteq S_a \times S_b$ *an expansion of* $Q_a \times Q_b$ if the following hold:

a. There exists a CPS $\mu_a \in \mathcal{C}(S_b)$ so that
   - (i) $Q_a \subseteq P_a = \rho_a(\mu_a)$,
   - (ii) μ_a strongly believes Q_b, and
   - (iii) if r_a is optimal under $\mu_a(\cdot|S_b)$ then $\pi_a(r_a, s_b) = \pi_a(s_a, s_b)$ for all $(s_a, s_b) \in Q_a \times Q_b$.

b. And, likewise, there is a CPS $\mu_b \in \mathcal{C}(S_a)$ satisfying analogous conditions.

Notice that we define only an expansion of a set $Q_a \times Q_b$ if $Q_a \times Q_b$ is a constant set. Also, if $P_a \times P_b$ is an expansion of $Q_a \times Q_b$, then there are CPS's μ_a and μ_b that satisfy conditions (i)–(iii) of Definition 21. We refer to these as *the associated CPS's*.

LEMMA 11. *Fix a PI game satisfying NRT. Suppose $P_a \times P_b$ is an expansion of $Q_a \times Q_b$, and fix associated CPS's μ_a and μ_b. Let X_a be the set of strategies that are optimal under $\mu_a(\cdot|S_b)$ and likewise define X_b. Then $X_a \times X_b$ is a constant set.*

PROOF. Since $P_a \times P_b$ is an expansion of $Q_a \times Q_b$, then $Q_a \times Q_b$ is a constant set. (This is by definition.) It follows from condition (iii) of Definition 21 and NRT that $X_a \times Q_b$ and $Q_a \times X_b$ are constant sets. Then using NRT, each profile in $X_a \times Q_b$ reaches the same terminal node. And likewise for $Q_a \times X_b$. In fact, the terminal node reached by $X_a \times Q_b$ and $Q_a \times X_b$ must be the same one, since $(X_a \times Q_b) \cap (Q_a \times X_b) = (Q_a \times Q_b)$. Now fix a profile $(s_a, r_b) \in (X_a \setminus Q_a) \times (X_b \setminus Q_b)$. Note that there is a profile $(s_a, s_b) \in (X_a \setminus Q_a) \times Q_b$ and a profile $(r_a, r_b) \in Q_a \times (X_b \setminus Q_b)$. These profiles reach the same terminal node and so (s_a, r_b) must also reach that terminal node. This establishes that $X_a \times X_b$ is a constant set. □

Corollary 4. *Fix a PI game satisfying NRT. If $P_a \times P_b$ is an expansion of some $Q_a \times Q_b$, then $P_a \times P_b$ is constant.*

The next result is standard, so the proof is omitted.

Lemma 12. *Fix a measure $\varpi_a \in \mathcal{P}(S_b)$ so that s_a is optimal under ϖ_a given S_a. Then, for any information set h with $s_a \in S_a(h)$ and $\varpi_a(S_b(h)) > 0$, s_a is optimal under $\varpi_a(\cdot|S_b(h))$ given $S_a(h)$.*

Lemma 13. *Fix a PI game that satisfies NRT. If $P_a \times P_b$ is an expansion of $Q_a \times Q_b$, then there exists some $W_a \times W_b$ that is an expansion of $P_a \times P_b$.*

Proof. Begin with the fact that $P_a \times P_b$ is an expansion of $Q_a \times Q_b$ and choose an associated CPS μ_a (resp. μ_b) that satisfies the conditions of Definition 21. Let X_a (resp. X_b) be the set of strategies that are optimal under $\mu_a(\cdot|S_b)$ (resp. $\mu_b(\cdot|S_a)$). By Lemma 11, $X_a \times X_b$ is a constant set.

Construct a measure $\varpi_a \in \mathcal{P}(S_b)$ as follows: Begin with a measure $\overline{\varpi}_a$ with $\operatorname{Supp} \overline{\varpi}_a = P_b$. Construct ϖ_a so that, for each $r_b \in P_b$,

$$\varpi_a(r_b) = (1-\varepsilon)\mu_a(r_b|S_b) + \varepsilon\overline{\varpi}_a(r_b),$$

where $\varepsilon \in (0,1)$. Note that μ_a strongly believes $Q_b \subseteq P_b$ so $\operatorname{Supp} \mu_a(\cdot|S_b) \subseteq P_b$. With this and the fact that $\operatorname{Supp} \overline{\varpi}_a = P_b$, we have $\operatorname{Supp} \varpi_a = P_b$. Using the fact that $X_a \times P_b$ is a constant set, then $\pi_a(s_a, \varpi_a) = \pi_a(r_a, \varpi_a)$ for all $s_a, r_a \in X_a$. Moreover, when ε is sufficiently small, $\pi_a(s_a, \varpi_a) > \pi_a(r_a, \varpi_a)$ for all $s_a \in X_a$ and $r_a \in S_a \setminus X_a$. So we can choose ϖ_a so that s_a is optimal under ϖ_a if and only if $s_a \in X_a$.

Now construct a CPS $\nu_a \in \mathcal{C}(S_b)$ as follows: If $P_b \cap S_b(h) \neq \varnothing$, let $\nu_a(\cdot|S_b(h)) = \varpi_a(\cdot|S_b(h))$. (This is well defined since, in this case, $\varpi_a(S_b(h)) > 0$.) If $P_b \cap S_b(h) = \varnothing$, let $\nu_a(\cdot|S_b(h)) = \mu_a(\cdot|S_b(h))$. Lemma 9 establishes that $\nu_a(\cdot|\cdot)$ is a CPS. Construct a measure $\varpi_b \in \mathcal{P}(S_a)$ and a CPS $\nu_b \in \mathcal{C}(S_a)$ analogously.

Take $W_a = \rho_a(\nu_a)$ and $W_b = \rho_b(\nu_b)$. We show that $W_a \times W_b$ is an expansion of $P_a \times P_b$.

Begin with condition (i). By definition, $W_a = \rho_a(\nu_a)$. So, we need to show only that $P_a \subseteq W_a$. Fix some $s_a \in P_a$. By construction, s_a is optimal under ϖ_a. Let $h \in H_a$ with $s_a \in S_a(h)$. If $P_b \cap S_b(h) \neq \varnothing$, then $\varpi_a(\cdot|S_b(h)) = \nu_a(\cdot|S_b(h))$ and s_a is optimal under $\nu_a(\cdot|S_b(h))$ among all strategies in $S_a(h)$. (See Lemma 12.) If $P_b \cap S_b(h) = \varnothing$, then $\nu_a(\cdot|S_b(h)) = \mu_a(\cdot|S_b(h))$. So, again, s_a is optimal under $\nu_a(\cdot|S_b(h))$ given all strategies in $S_a(h)$. With this, $s_a \in \rho_a(\nu_a(\cdot|\cdot))$, as required.

Next, turn to condition (ii). We need to show that ν_a strongly believes P_b. For this, notice that if $P_b \cap S_b(h) \neq \varnothing$, then $\nu_a(P_b|S_b(h)) = \varpi_a(P_b|S_b(h)) = 1$.

Finally, we show condition (iii). Suppose r_a is optimal under $\nu_a(\cdot|S_b)$. We show that $\pi_a(r_a, s_b) = \pi_a(s_a, s_b)$ for all $(s_a, s_b) \in P_a \times P_b$. To see this, recall, $\nu_a(\cdot|S_b) = \varpi_a$. So if r_a is optimal under $\nu_a(\cdot|S_b)$, then $r_a \in X_a$. The claim now follows from the fact that $X_a \times X_b$ is a constant set that contains $P_a \times P_b$.

Replacing b with a establishes that $W_a \times W_b$ is an expansion of $P_a \times P_b$. □

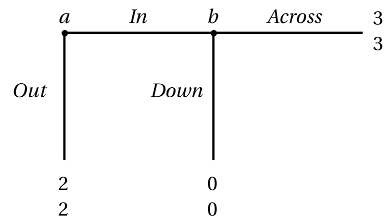

FIGURE 11. A PI game with NRT.

LEMMA 14. *Fix a PI game that satisfies NRT. Let (s_a, s_b) be a Nash equilibrium in sequentially justifiable strategies. Then there exists an EFBRS, viz. $Q_a \times Q_b$, that contains (s_a, s_b).*

PROOF. Fix a Nash equilibrium in sequentially optimal strategies, viz. (s_a, s_b). Let $Q_a^0 \times Q_b^0 = \{s_a\} \times \{s_b\}$. By Lemma 10, $Q_a^0 \times Q_b^0$ satisfies the best response property. So there is a CPS μ_a (resp. μ_b) that strongly believes $\{s_b\}$ (resp. $\{s_a\}$) and so that s_a (resp. s_b) is sequentially optimal under μ_a (resp. μ_b). Let $Q_a^1 = \rho_a(\mu_a)$ (resp. $Q_b^1 = \rho_b(\mu_a)$). Note that $Q_a^1 \times Q_b^1$ is an expansion of $Q_a^0 \times Q_b^0$ (associated with the CPS's μ_a and μ_b). Now repeatedly apply Lemma 13 to get sets $Q_a^0 \times Q_b^0, Q_a^1 \times Q_b^1, Q_a^2 \times Q_b^2, \ldots$, where each $Q_a^{m+1} \times Q_b^{m+1}$ is an expansion of $Q_a^m \times Q_b^m$. Since the game is finite, there is some M with $Q_a^m \times Q_b^m = Q_a^M \times Q_b^M$ for all $m \geq M$. The set $Q_a^M \times Q_b^M$ is an EFBRS. □

D.IV *Closing the gap*

In the text, we mentioned that there is a gap between parts (i) and (ii) of Proposition 3.

We begin by pointing out that we cannot improve part (ii) to say that, starting from any pure Nash equilibrium, we get an EFBRS. To see this, refer to Figure 11. There is a unique EFBRS, namely $\{In\} \times \{Across\}$. That said, the pair $(Out, Down)$ is a Nash equilibrium—of course, it is not a Nash equilibrium in sequentially justifiable strategies.

We do not know if part (i) can be improved to read, If $Q_a \times Q_b$ satisfies the best response property, then each $(s_a, s_b) \in Q_a \times Q_b$ is outcome equivalent to a sequentially justifiable Nash equilibrium. Let us better understand the problem.

Return to Lemma 7 and the proof thereof. Suppose, we strengthened the induction hypothesis so that we can look at a sequentially justifiable Nash equilibrium of subgame 1, viz. (r_a^1, r_b^1). Following the proof, we use this to construct a Nash equilibrium $(r_a, (r_b^1, \underline{q}_b^2, \ldots, \underline{q}_b^K))$, where each $\underline{q}_b^k$ is the minimax strategy on subtree k. But now we need to show that the constructed equilibrium is sequentially justifiable. Here is where the problem arises: the strategy $\underline{q}_b^k$ (on subtree k) may not be a best response to any strategy on that subtree. Thus, the proof breaks down. Of course, it may very well be that there is another method of proof.

In the text, we mentioned a related result (Proposition 4), which speaks to the gap. To show this result, it suffices to show the following lemma.

LEMMA 15. *Suppose $Q_a \times Q_b$ is a constant set that satisfies the best response property. Then there exists a mixed-strategy Nash equilibrium, viz. (σ_a, σ_b), so that* *(i)* $Q_a \times Q_b$ *is outcome equivalent to* (σ_a, σ_b) *and*

(ii) each $s_a \in \operatorname{Supp} \sigma_a$ *(resp.* $s_b \in \operatorname{Supp} \sigma_b$*) is sequentially justifiable.*

Proof. Pick some $(r_a, r_b) \in Q_a \times Q_b$ and let $\mu_a \in \mathcal{C}(S_b)$ be a CPS so that $r_a \in \rho_a(\mu_a)$ and μ_a strongly believes Q_b. Set $\sigma_b = \mu_a(\cdot|S_b)$. Construct σ_a analogously.

First notice that (σ_a, σ_b) is a mixed-strategy Nash equilibrium. Begin by using the fact that $\mu_b(Q_a|S_a) = 1$ and $\mu_a(Q_b|S_b) = 1$. As such, $\operatorname{Supp} \sigma_a \times \operatorname{Supp} \sigma_b \subseteq Q_a \times Q_b$. Since $Q_a \times Q_b$ is a constant set, for each $(s_a, s_b) \in \operatorname{Supp} \sigma_a \times \operatorname{Supp} \sigma_b$, $\pi(s_a, s_b) = \pi(r_a, r_b)$. So for each $s_a \in \operatorname{Supp} \sigma_a$ and each $q_a \in S_a$,

$$\begin{aligned}\pi_a(s_a, \sigma_b) &= \pi_a(r_a, r_b) \\ &= \pi_a(r_a, \sigma_b) \geq \pi_a(q_a, \sigma_b),\end{aligned}$$

where the inequality holds because $r_a \in \rho_a(\mu_a)$ and $\mu_a(\cdot|S_b) = \sigma_b$. Applying an analogous argument to b establishes that (σ_a, σ_b) is indeed a Nash equilibrium.

Next notice that $Q_a \times Q_b$ is outcome equivalent to (σ_a, σ_b). To see this, recall that $\operatorname{Supp} \sigma_a \times \operatorname{Supp} \sigma_b \subseteq Q_a \times Q_b$ and $Q_a \times Q_b$ is a constant set. So it is immediate that, for each $(s_a, s_b) \in Q_a \times Q_b$, $\pi(s_a, s_b) = \pi(\sigma_a, \sigma_b)$.

Last, notice that each $s_a \in \operatorname{Supp} \sigma_a$ is sequentially justifiable and likewise for b. To see this, recall that $\operatorname{Supp} \sigma_a \times \operatorname{Supp} \sigma_b \subseteq Q_a \times Q_b$. So if $s_a \in \operatorname{Supp} \sigma_a$, then $s_a \in Q_a$ and so s_a is sequentially justifiable. □

The proof of Proposition 4 is immediate from Lemmata 8 and 15.

References

Battigalli, Pierpaolo (1996), "Strategic independence and perfect Bayesian equilibria." *Journal of Economic Theory*, 70, 201–234. [80]

Battigalli, Pierpaolo (1997), "On rationalizability in extensive form games." *Journal of Economic Theory*, 74, 40–61. [69, 74, 86]

Battigalli, Pierpaolo (1999), "Rationalizability in incomplete information games." Working Paper ECO 99/17, EUI. [69]

Battigalli, Pierpaolo and Marciano Siniscalchi (1999a), "Hierarchies of conditional beliefs and interactive epistemology in dynamic games." *Journal of Economic Theory*, 88, 188–230. [62, 69, 79]

Battigalli, Pierpaolo and Marciano Siniscalchi (1999b), "Interactive beliefs, epistemic independence and strong rationalizability." *Research in Economics*, 53, 247–273. [81]

Battigalli, Pierpaolo and Marciano Siniscalchi (2002), "Strong belief and forward induction reasoning." *Journal of Economic Theory*, 106, 356–391. [57, 58, 60, 62, 69, 79]

Battigalli, Pierpaolo and Marciano Siniscalchi (2003), "Rationalization and incomplete information." *Advances in Theoretical Economics*, 3 (1). [57, 58, 69, 70, 80]

Battigalli, Pierpaolo and Marciano Siniscalchi (2007), "Interactive epistemology in games with payoff uncertainty." *Research in Economics*, 61, 165–184. [80]

Battigalli, Pierpaolo and Amanda Friedenberg (2009), "Context-dependent forward induction reasoning." Working Paper 351, IGIER, Università Bocconi. [61, 79, 82, 86]

Battigalli, Pierpaolo and Andrea Prestipino (2011), "Transparent restrictions on beliefs and forward induction reasoning in games with asymmetric information." Working Paper 376, IGIER, Università Bocconi. [70, 80]

Ben-Porath, Elchanan (1997), "Rationality, Nash equilibrium and backwards induction in perfect-information games." *Review of Economic Studies*, 64, 23–46. [77, 92]

Ben-Porath, Elchanan and Eddie Dekel (1992), "Signaling future actions and the potential for sacrifice." *Journal of Economic Theory*, 57, 36–51. [81]

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, London. [62]

Brandenburger, Adam, Amanda Friedenberg, and H. Jerome Keisler (2008), "Admissibility in games." *Econometrica*, 76, 307–352. [72]

Brandenburger, Adam and Amanda Friedenberg (2010), "Self-admissible sets." *Journal of Economic Theory*, 145, 785–811. [77, 78, 85, 91]

Fudenberg, Drew and Jean Tirole (1991), "Perfect Bayesian equilibrium and sequential equilibrium." *Journal of Economic Theory*, 53, 236–260. [80]

Hammond, Peter J. (1987), "Extended probabilities for decision theory and games." Unpublished paper, Department of Economics, Stanford University. [80]

Marx, Leslie and Jeroen Swinkels (1997), "Order independence for iterated weak dominance." *Games and Economic Behavior*, 18, 219–245. [77]

Osborne, Martin J. and Ariel Rubinstein (1994), *A Course in Game Theory*. MIT Press, Cambridge, Massachusetts. [58, 59]

Pearce, David G. (1984), "Rationalizable strategic behavior and the problem of perfection." *Econometrica*, 52, 1029–1050. [58]

Reny, Philip J. (1993), "Common belief and the theory of games with perfect information." *Journal of Economic Theory*, 59, 257–274. [75]

Rényi, Alfréd (1955), "On a new axiomatic theory of probability." *Acta Mathematica Hungarica*, 6, 285–335. [60]

Stalnaker, Robert (1998), "Belief revision in games: Forward and backward induction." *Mathematical Social Sciences*, 36, 31–56. [58]

Submitted 2009-7-30. Final version accepted 2010-12-4. Available online 2010-12-7.