



Institutional Members: CEPR, NBER and Università Bocconi

WORKING PAPER SERIES

Belief change, Rationality, and Strategic Reasoning in Sequential Games

Pierpaolo BATTIGALLI, Emiliano CATONINI, Julien MANILI

Working Paper n. 679

This Version: August 22nd, 2022

IGIER – Università Bocconi, Via Guglielmo Röntgen 1, 20136 Milano –Italy
<http://www.igier.unibocconi.it>

The opinions expressed in the working papers are those of the authors alone, and not those of the Institute, which takes non institutional policy position, nor those of CEPR, NBER or Università Bocconi.

Belief change, Rationality, and Strategic Reasoning in Sequential Games*

P. Battigalli,* E. Catonini,[◦] J. Manili[†]

* Bocconi University and IGIER

[◦] NYU Shanghai

[†] Northwestern University

This (extended) version: August 22, 2022. First version: May 15, 2021.

Abstract

A central aspect of strategic reasoning in sequential games consists in anticipating how co-players would react to information about past play, which in turn depends on how co-players update and revise their beliefs. Several notions of belief system have been used to model how players' beliefs change as they obtain new information, some imposing considerably more discipline than others on how beliefs at different information sets are related. We highlight the differences between these notions of belief system in terms of introspection about one's own conditional beliefs, but we also show that such differences do not affect the essential aspects of rational planning and the behavioral implications of strategic reasoning, as captured by rationalizability.

KEYWORDS: Sequential games, chain rule, partial introspection, rational planning, rationalizability.

JEL Codes: C72, C73, D83.

1 Introduction

In sequential games, the beliefs of players may change as the play unfolds. In particular, a player may find out that her co-players are implementing strategies to which she initially assigned zero probability. In this case, the player has to form a new belief, which has to be consistent with her information about the behavior of the other players. We call the latter

*We thank Nicola Bariletto, Nicodemo De Vito, Amanda Friedenberg, Pierfrancesco Guarino, Anna Merotto, and Marciano Siniscalchi for useful comments.

process belief **revision**, which contrasts with the mere **updating** of beliefs after observing events that were assigned strictly positive probability.

Several notions of belief system have been used to model how players' beliefs change as they receive new information, ranging from "updating *previous* beliefs whenever possible" (cf. Ben-Porath 1997, Battigalli & Siniscalchi 2003), which we call **forward consistency**, to notions incorporating considerable discipline on how players revise their beliefs in relation to the beliefs they would hold at counterfactual contingencies. The most restrictive one, **complete consistency**, obtains from considering complete conditional probability systems (Myerson 1986, Battigalli 1996), that is, by requiring that players update, or revise their beliefs *as if* they could hypothetically condition on any nonempty event about co-players' behavior in compliance with the chain rule. An intermediate notion of "**standard**" consistency obtains from conditional probability systems over the observable events about co-players' behavior induced by information sets (cf. Renyi 1955, Battigalli & Siniscalchi 2002).

Yet the implications of these requirements for optimal planning and strategic reasoning remain unclear. A central aspect of strategic reasoning in sequential games consists in anticipating how co-players would react to expected and unexpected information about past play, which in turn depends on how co-players update and revise their beliefs. So the restrictions incorporated by these models of belief change may affect game-theoretic predictions.

To tackle this issue, we first shed light on the three nested notions of belief consistency we introduced above, by interpreting them in terms of different degrees of introspection about one's own conditional beliefs, which can be naturally associated with different notions of rational planning. We observe that forward planning—as captured by weak sequential optimality of strategies—does not require a player to anticipate her own beliefs at unexpected contingencies. Forward consistency obtains as an application of the chain rule, precisely for players who do not ask themselves what beliefs they would hold at information sets that they currently deem impossible. Instead, the most demanding notion of folding-back planning—characterized by sequential optimality—requires full introspection, that is, a player's ability to anticipate her beliefs at all information sets.¹ Then, additional restrictions, such as "same information about others implies same beliefs," naturally emerge as a result of the cognitive rationality of introspective players. This gives rise to standard consistency, which requires beliefs about other players to depend only on information about others' behavior and not on own past behavior. Complete consistency implies further discipline on the relative probabili-

¹Traditionally, folding-back planning is presented as an algorithm to compute optimal plans in *decision trees*, starting from terminal decision nodes and working backward. Traditional decision trees do not feature decision situations that the decision maker deems impossible. This is different from the tree of information sets of a given player in a game, whose occurrence depend on co-players' behavior that *ex ante* may be deemed impossible. In Section 6, we sketch a weaker notion of *partial* folding-back planning performed on the restricted tree of information sets a player deems possible, starting all over again when surprised. This only requires partial introspection.

ties of strategies at different information sets, even if such information sets do not correspond to nested information about others' behavior. Such discipline is implied by a player's introspective ability to determine what she would believe if she *hypothetically* knew that the other players were implementing strategies within an arbitrary subset, even if it does not correspond to an information set.

In light of these differences, the following question arises: To what extent do the behavioral implications of rationality and strategic reasoning depend on the degree of belief consistency? Or, given the above interpretation, does assuming that players are fully introspective or only partially introspective make a difference for the analysis? Perhaps surprisingly, we find that the degree of consistency in belief revision is irrelevant for the strategic analysis in the following sense: the foregoing three notions of consistent belief change yield equivalent versions of rationalizability for sequential games. We illustrate the main insights concerning this invariance result with our running example (see Figure 1). We focus on strong rationalizability (Pearce 1984, Battigalli 1997), a procedure that captures forward-induction reasoning. In Section 8 we argue that the invariance result extends to other notions of rationalizability such as backwards rationalizability (Perea 2014, Battigalli & De Vito 2021) and initial rationalizability (Ben-Porath 1997, Battigalli & Siniscalchi 1999), but we also point out that solution concepts featuring contextual restrictions on belief systems (Battigalli & Siniscalchi 2003, Battigalli & Friedenberg 2013) may be affected by the assumed notion of consistency.




Figure 1: Γ' , a common interest game between Isa and Joe.

Heuristic analysis of an example Consider the game depicted in Figure 1. At the beginning of the game, Isa (i) chooses between Left and Right, and Joe (j) simultaneously chooses between Quit (which terminates the game) and Continue. If Joe Continues, Isa observes this and decides whether to go across (terminating the game), or down. If Isa goes down, Joe observes this but not her initial move, and chooses between left and right. Suppose that Isa is initially certain that Joe would Quit. Then she would have to revise her beliefs after observing either (L, C) or (R, C) . Can Isa form different beliefs about Joe after (L, C) and (R, C) , although they reveal the *same information* about Joe’s behavior? According to forward consistency, she can, because (L, C) and (R, C) do not precede each other. Thus, forward consistency allows beliefs about *others* to depend on *own* actions. This is ruled out by standard and complete consistency, whereby same information about co-players implies same beliefs.

Can these differences matter for the analysis of the behavioral implications of rationality and strategic reasoning? Our results provide a negative answer. Intuitively, the reason is that weak sequential optimality of a strategy, which characterizes the observable behavioral implications of rationality, depends only on the beliefs at the information sets that are not precluded by the strategy itself. For instance, suppose that Isa is initially certain of Quit. In this case, how she would revise her belief about Joe upon observing Continue is irrelevant for the subjective optimality of her first action. If Isa finds it optimal to go Right and Joe unexpectedly Continues, Isa’s counterfactual belief had she played Left does not matter to determine the optimality of her next move at (R, C) , because only her belief given (R, C) matters. Thus, whether the belief she would have held after (L, C) is equal to the belief after (R, C) —as implied by standard and complete consistency—is inconsequential. We will explain this in detail in our analysis of rational planning and justifiable behavior of Sections 6 and 7.

The conceptual implications of these considerations are better understood by considering our interpretation of forward consistency and how it relates to rational planning. We provide here an informal explanation of why Isa might have different beliefs about Joe conditional on (L, C) and (R, C) , despite the fact that at both nodes she has the same information about Joe. As anticipated, according to subjective expected utility maximization, in order to decide what to do at the root, Isa only has to consult her initial belief. If she is certain of Quit, then Right is the best choice (one of the best, if $\varepsilon = 0$). She has no need for further planning and she comes up with the *partial* strategy of going Right. Thus, she may be only *partially introspective and unable to anticipate how she would revise her belief if surprised*. Such partial introspection may prevent her from imposing the cognitive rationality rule “same information about others implies same belief,” unless such rule is in somehow wired into her belief formation process. Of course, if surprised, Isa would form *some* revised belief and

make a choice based on it. For example, if she initially goes Right and is surprised by C , it may be the case that she is pre-disposed to believe that Joe is more likely to continue with right, and thus choose to go down. In this case, she would implement the reduced strategy $R.d''$.² But this does not imply that she had initially planned to go down if surprised. In other words, $R.d''$ is just a *description* of Isa’s behavior, but not necessarily a plan of Isa. What matters for strategic reasoning is to be able to anticipate the behavior of others. From this perspective, Joe’s belief in the rationality of Isa allows for the possibility that he assigns positive probability to her behavior being described by $R.d''$. Thus, our results about the invariance of rationalizability to cognitive rationality restrictions on belief systems beyond forward consistency may be interpreted as saying that the underlying epistemic justifications of rationalizability do not rely on full introspection.

The rest of the paper is organized as follows: Section 2 introduces the game-theoretic set-up and the notion of reduced strategy. Section 3 presents our notions of conditional belief systems and a key lemma about their relationships. Section 4 characterizes standard belief systems in games with observable actions, and shows that standard and forward consistent belief systems coincide in games with perfect information. Section 5 characterizes complete consistency. Section 6 presents notions of rational planning along with our interpretation based on partial or full introspection about own beliefs. Section 7 presents strong rationalizability and uses the key lemma of Section 3 to show that its predictions are invariant to the assumed notion of belief change. Section 8 discusses extensions (including other notions of rationalizability and sets with the best response property) and the related literature. Section 9 concludes. Appendix A collects some proofs omitted from the main text, Appendix B is devoted to backwards rationalizability, and Appendix C carries out a comparison between complete consistency and the notion of consistency of Kreps and Wilson (1982).

2 Sequential games with perfect recall

Our analysis is restricted to *finite* sequential games without chance moves played by agents with *perfect recall*, represented in extensive form.³ Some knowledge of the extensive and

²Rubinstein (1991) calls such reduced strategies “plans of action.” See also Osborne & Rubinstein (1994, p. 103). Reduced strategies should not be confused with partial strategies, such as going Right in the example, which have to be extended upon observing unexpected co-players’ behavior.

³Games where—according to the rules—players are always reminded of the information previously provided to them and of the actions they took are extremely rare. For this reason, we interpret the perfect recall property as resulting from the interaction between players’ personal cognitive features and the rules of the game. When players’ mnemonic abilities are perfect, whatever the rules of the game, their information structure may be represented with information partitions satisfying perfect recall (see Battigalli & Generoso 2021). Also, we refrain from using expressions like “extensive-form game” or “normal-form game”, because by “game” we mean the object being represented, not its mathematical representation. With this, the extensive

strategic-form representations of sequential games is taken for granted. Thus, for the primitive terms of the analysis, only the necessary symbols and definitions with rather terse explanations are given below. The reader interested in the details should consult, e.g., Selten (1975), or Osborne & Rubinstein (1994). We instead expand on the interpretation of some derived terms. Note also that the formalism used here is more expressive than the traditional one due to Kuhn (1953), because (i) it represents simultaneous moves directly by letting plays be sequences of action profiles, and (ii) it represents also the information of inactive players, which is potentially relevant for our analysis of belief change.

A **sequential game** (played by agents) **with perfect recall** is a structure

$$\Gamma = \langle I, \bar{X}, (A_i, X_i, H_i, u_i)_{i \in I} \rangle$$

where I is a finite set of **players**, A_i is a nonempty finite set of potentially available **actions** for player i , and \bar{X} is a *finite tree* of feasible sequences of action profiles, called **histories** or **nodes**.⁴ We let \emptyset denote the **empty sequence** (root of tree \bar{X}), Z denote the set of **terminal** histories (or **paths**), and $X = \bar{X} \setminus Z$ denote the set of **non-terminal** histories. We write $x \prec x'$ ($x \preceq x'$) if sequence x is a strict (weak) prefix of x' , that is, node x precedes node x' in tree \bar{X} . For each player $i \in I$ there is a subset $X_i \subseteq X$ of nonterminal nodes where i is **alert**, i.e., she processes information; $\iota(x) = \{i \in I : x \in X_i\}$ denotes the set of players who are alert at x and we assume that $\iota(\emptyset) = I$, that is, all players are alert at the root. Alert players are **active** if they can choose between two or more alternative actions, and **inactive** otherwise, i.e., if all they can do is to “wait.” This is described by a profile of nonempty-valued feasibility correspondences $(\mathcal{A}_i(\cdot) : X_i \rightrightarrows A_i)_{i \in I}$ such that, for every $x \in X$, $(x, a_{\iota(x)}) \in \bar{X}$ if and only if $a_{\iota(x)} \in \times_{i \in \iota(x)} \mathcal{A}_i(x)$.⁵ Thus, $i \in \iota(x)$ is inactive at x if $\mathcal{A}_i(x)$ is a singleton. To avoid trivialities, we assume that, for each $x \in X$, *at least one of the alert players* (those in $\iota(x)$) *is active*. In our graphical representations, such as the game tree Γ' in Figure 1, we only show the actions and information of active players. The information structure of player i is represented by the collection H_i of **information sets** of player i , where:

1. H_i is a *partition* of X_i ;
2. for every $h_i \in H_i$ and $x, x' \in h_i$, $\mathcal{A}_i(x) = \mathcal{A}_i(x') =: \mathcal{A}_i(h_i)$;
3. (*perfect recall*) for every $h_i \in H_i$ and $x, x' \in h_i$ with $x \neq x'$, we have (i) $x \not\prec x'$ and

and normal forms are types of representation, not types of game.

⁴An action profile of a nonempty subset $J \subseteq I$ of players is an element of the cross-product $\times_{i \in J} A_i$. A set of sequences is a **tree** if it contains every prefix of each one of its elements, including the empty sequence.

⁵This property ensures that what a player can do at a node does not depend on what co-players are simultaneously doing—otherwise actions would not be simultaneous.

(ii) for all $(\tilde{x}, a) \preceq x$ with $\tilde{x} \in \tilde{h}_i$ for some $\tilde{h}_i \in H_i$, there exists $(\tilde{x}', a') \preceq x'$ such that $\tilde{x}' \in \tilde{h}_i$ and $a_i = a'_i$.⁶

Perfect recall implies that we can unambiguously partially order the information sets of each player i with the “prefix of” precedence relation of X : for all distinct $h_i, h'_i \in H_i$, h_i is a **predecessor** of h'_i , denoted $h_i \prec_i h'_i$, if and only if every history in h'_i follows some history in h_i (i.e., it has a prefix in h_i). With this, (H_i, \prec_i) is a tree.⁷

Finally, $u_i : Z \rightarrow \mathbb{R}$ is the **payoff function** of player i . For the sake of simplicity, many of our examples feature **common interests** (CI): $u_i = u_j$ for all $i, j \in I$.

We consider two notable special cases: a game has **observable actions** (or perfect monitoring) if $X_i = X$ for every i (players are always alert) and all information sets are singletons, in which case we write $H_i = X =: H$ for every i and we do not distinguish between histories/nodes and the singleton information sets containing them. A game with observable actions has **perfect information** if, for every history $h \in H$, only one player is active. For example, game Γ'' depicted in Figure 2 has observable actions, whereas game Γ''' depicted in Figure 3 has perfect information.




Figure 2: Γ'' , a CI game with observable actions.

A **strategy** for player i is a function $s_i : H_i \rightarrow A_i$ that assigns to each information set $h_i \in H_i$ a feasible action $s_i(h_i) \in \mathcal{A}_i(h_i)$. Thus, the set of strategies of player i is the cross-product of feasible action sets, $S_i = \times_{h_i \in H_i} \mathcal{A}_i(h_i)$. We denote by $S = \times_{i \in I} S_i$ the set

⁶Property (i) says that players cannot end up twice in the same information set because they remember having moved before. Property (ii) says that if two histories are in the same information set, then player i must have been unable to distinguish the prefixes of these histories at earlier information sets and must have taken the same actions at such earlier information sets (since she recalls her past information and actions).

⁷Battigalli & Generoso (2021) show that information sets can be associated with personal histories of own actions and messages.




Figure 3: Γ''' , a CI game with perfect information.

of strategy profiles and by $S_{-i} = \times_{j \neq i} S_j$ the set of i 's co-players' strategy profiles. The implementation of a profile of strategies $s \in S$ induces a unique terminal history $\zeta(s)$, where $\zeta : S \rightarrow Z$ denotes the **path function**. Although we do not require feasible action sets at distinct information sets of an active player to be disjoint, this condition holds in our examples. This eases notation, allowing us to write the strategies of our examples as lists of actions separated by dots, such as $R.a'.d''$ for Isa in game Γ' of Figure 1.

For each $h_i \in H_i$, the set of strategy profiles compatible with information set h_i is

$$S(h_i) = \{s \in S : \exists x \in h_i, x \prec \zeta(s)\}.$$

Let $S_i(h_i) = \text{proj}_{S_i} S(h_i)$ and $S_{-i}(h_i) = \text{proj}_{S_{-i}} S(h_i)$. It is worth noting that perfect recall implies the following: for all $h_i, h'_i \in H_i$, we have (i) $S(h_i) = S_i(h_i) \times S_{-i}(h_i)$, (ii) if h'_i follows h_i , then $S(h'_i) \subseteq S(h_i)$, hence, $S_{-i}(h'_i) \subseteq S_{-i}(h_i)$, (iii) $S(h_i) \cap S(h'_i) \neq \emptyset$ if and only if either $h_i \preceq_i h'_i$ or $h'_i \preceq_i h_i$.⁸ Yet, it is possible that $S_{-i}(h_i) \cap S_{-i}(h'_i) \neq \emptyset$ even if h_i and h'_i are not ordered, because the information sets of i also encode information about her past actions. For example, in game Γ' , $S_{-i}(\{(L, C)\}) = \{C.l, C.r\} = S_{-i}(\{(R, C)\})$.

As in most of the work on strategic reasoning, here *strategies represent both contingent plans in the minds of rational players and descriptions of information-dependent behavior*. Thus, as a player plans her strategy s_i , she assesses the likelihood of the possible “ways of behaving,” or “action rules” of the others, s_{-i} .

The interpretation of strategies as plans or mere descriptions of behavior is related to an important structural equivalence relation. Let $H_i(s_i) = \{h_i \in H_i : s_i \in S_i(h_i)\}$ denote the collection of information sets that may occur if i plays strategy s_i . For example, in game Γ' , $H_i(R.a'.a'') = H_i(R.a'.d'') = \{\{\emptyset\}, \{(R, C)\}\}$; in game Γ'' , $H_j(Q.x.y) = \{\{\emptyset\}\}$ and $H_j(C.x.y) = \{\{\emptyset\}, \{(L, C)\}, \{(R, C)\}\}$ for all $(x, y) \in \{\ell', r'\} \times \{\ell'', r''\}$.

⁸To see this, if $s \in S(h_i)$ and $s \in S(h'_i)$, then there exist $x \in h_i$ such that $x \prec \zeta(s)$ and $x' \in h'_i$ such that $x' \prec \zeta(s)$. Hence, either $x' \preceq x$ or $x \preceq x'$, which implies by perfect recall that either $h_i \preceq_i h'_i$ or $h'_i \preceq_i h_i$. The reason is that perfect recall implies that for any $h_i, h'_i \in H_i$, if $x \in h_i$, $x' \in h'_i$ and $x \prec x'$, then all nodes of h'_i are preceded by some node of h_i , that is, $h_i \prec_i h'_i$.

Definition 1. Two strategies $s'_i, s''_i \in S_i$ are (1) **behaviorally equivalent** if $H_i(s'_i) = H_i(s''_i)$ and $s'_i(h_i) = s''_i(h_i)$ for every $h_i \in H_i(s'_i)$, (2) **realization-equivalent** if $\zeta(s'_i, s_{-i}) = \zeta(s''_i, s_{-i})$ for every $s_{-i} \in S_{-i}$.

Kuhn (1953) proved that these two equivalence relations coincide:

Remark 1. (Kuhn, 1953, Theorem 1) Two strategies are behaviorally equivalent if and only if they are realization-equivalent.

Let \equiv_i denote this (behavioral or realization) equivalence relation. We call the elements of the quotient set $S_i | \equiv_i$ “structurally reduced strategies,” abbreviated in “**reduced strategies**”; e.g., Isa has $2^3 = 8$ strategies, but only 4 reduced strategies in Γ' : $L.a'$, $L.d'$, $R.a''$, and $R.d''$, with $L.a' = \{L.a'.a'', L.a'.d''\}$ etc. Often these equivalence classes are instead called “plans of action,” suggesting that how s_i is defined outside of $H_i(s_i)$ is irrelevant for planning, and sometimes adding that the restriction of s_i to $H_i \setminus H_i(s_i)$ should be interpreted as an expectation of the co-players (e.g., Osborne & Rubinstein 1994, p. 103). We do not adopt this terminology because we have a different perspective that will be fully spelled out in Section 6.

To anticipate, one may think of a player *planning forward* or *backward*. Planning forward means comparing the expected payoffs of different courses of action like Q (quit), or $C.x.y$ (continue, then x if L and y if R) for Joe in game Γ'' (opting for $C.\ell'.r''$ if he initially assigns probability larger than 0.5 to L , and for Q otherwise). When planning forward, there is no need to specify actions for contingencies that cannot occur given the plan under consideration (in Γ'' , if Joe considers quitting, he does not have to plan what to do if he instead continues); thus, *reduced strategies correspond to forward plans*. Planning backward means first asking oneself what to do at any last move, using the answer as a contingent prediction about own behavior, and—given this—recursively planning what to do at earlier moves. For example, in game Γ'' , without having made up his mind about the move at the root, Joe first plans to choose ℓ' after (L, C) and r'' after (R, C) ; next, he plans to continue (quit) at the root if he initially assigns probability larger (smaller) than 0.5 to L . Whatever Joe plans to do at the root, *the result of such backward planning is a complete strategy* $s_j = a_j.\ell'.r''$ with $a_j \in \{C, Q\}$.⁹

The foregoing arguments suggest that, if we regard relation \equiv_i as behavioral equivalence, we are led to interpret reduced strategies as “forward plans.” If we instead regard \equiv_i as realization equivalence, we can think of reduced strategies as *sufficient descriptions of i 's*

⁹It also makes sense to plan only for moves that are deemed possible under current beliefs. For example, if in game Γ' Isa is initially certain that Joe is going to quit, then she just plans what to do at the root, which yields a partial reduced strategy. If Joe continues, thus surprising Isa, she decides on the spot what to do next. See Section 6.

behavior in the eyes of the co-players (or an external observer). Indeed, if $s'_i \equiv_i s''_i$, independently of the co-players' behavior, it is impossible to distinguish between s'_i and s''_i by observing the realized path; furthermore, it is not necessary for i 's co-players to distinguish between s'_i and s''_i in order to assess the likely consequences of taking different actions at (or implementing different continuation strategies starting from) an information set. For example, all that matters for Joe in game Γ' of Figure 1 are the probabilities of the reduced strategies $L.a'$, $L.d'$, $R.a''$, and $R.d''$, interpreted as sufficient descriptions of Isa's behavior.

3 Conditional beliefs

In this section, we introduce several representations of beliefs for sequential games. Players are uncertain about how the others would behave in various circumstances, and form beliefs on others' behavior to assess the likely consequences of their actions. The beliefs of a player will typically change as the game progresses. At information set $h_i \in H_i$, player i learns that the co-players are behaving according to a strategy profile in $S_{-i}(h_i)$. If player i 's beliefs before the realization of information set h_i assigned probability 0 to $S_{-i}(h_i)$, then the realization of h_i falsifies i 's earlier beliefs, which have to be revised, rather than just updated according to the rules of conditional probability. We thus have to model what players would believe in all circumstances, and how their beliefs change as they receive new information.

We first consider an abstract representation of conditional thinking by means of “conditional probability systems” (Renyi 1955), which requires some consistency between beliefs conditional on different events. Next we move to “systems of beliefs,” which specify the beliefs a player would hold at each information set. We introduce different degrees of consistency among beliefs at different information sets and relate them to conditional probability systems.

We use the following notation: For any finite set Ω , interpreted as the space of uncertainty, let $\Delta(\Omega)$ be the set of probability measures on Ω . Whenever the underlying space of uncertainty Ω is understood, for all events $E \subseteq \Omega$ we let $\Delta(E) = \{\mu \in \Delta(\Omega) : \mu(E) = 1\}$.

3.1 Conditional probability systems

Let Ω be a finite space of uncertainty. Fix a nonempty collection of “conceivable conditioning events” $\mathcal{C} \subseteq 2^\Omega \setminus \{\emptyset\}$. We call the pair (Ω, \mathcal{C}) a **conditional space**.¹⁰ For example, if we consider player i in a game, we have $\Omega = S_{-i}$, and a natural collection of conditioning events

¹⁰If Ω is infinite, consider a triple $(\Omega, \mathcal{B}, \mathcal{C})$, where \mathcal{B} is the relevant sigma-algebra and $\mathcal{C} \subseteq \mathcal{B}$.

is given by the observable events about co-players' behavior

$$\mathcal{H}_i := \{S_{-i}(h_i) \subseteq S_{-i} : h_i \in H_i\}.$$

Definition 2. Fix a conditional space (Ω, \mathcal{C}) . An array of probability measures $\mu = (\mu(\cdot|C))_{C \in \mathcal{C}} \in \times_{C \in \mathcal{C}} \Delta(C)$ is a **conditional probability system (CPS)** on (Ω, \mathcal{C}) , written $\mu \in \Delta^{\mathcal{C}}(\Omega)$, if it satisfies the *chain rule*: for all $E \subseteq \Omega$, $C, D \in \mathcal{C}$,

$$E \subseteq D \subseteq C \Rightarrow \mu(E|C) = \mu(E|D) \mu(D|C).$$

The CPSs on $(\Omega, 2^{\Omega} \setminus \{\emptyset\})$, whose set is denoted by $\Delta^*(\Omega)$, are called **complete**.

In words, whenever possible, the beliefs for distinct conditioning events must be related to each other by the standard rules of conditional probability. Indeed, the chain rule is equivalent to requiring that, for all $E \subseteq \Omega$ and $C, D \in \mathcal{C}$ with $E \subseteq D \subseteq C$,

$$\mu(D|C) > 0 \Rightarrow \mu(E|D) = \frac{\mu(E|C)}{\mu(D|C)}.$$

(Note that, since $0 \leq \mu(E|C) \leq \mu(D|C)$, $\mu(D|C) = 0$ implies that the chain-rule equality holds trivially as $0 = 0$.)

As anticipated, a natural representation of conditional beliefs in games is to consider CPSs on (S_{-i}, \mathcal{H}_i) , where player i forms beliefs conditional on each observable event about the behavior of others (cf. Battigalli & Siniscalchi 2002). Complete CPSs instead assume that players form their beliefs conditional on every possible event: for every C , the player asks herself “What would I believe if I hypothetically knew C ?” Event C need not represent information player i may obtain during the game, yet she can still ask herself this question and answer it with a belief $\mu^i(\cdot|C)$. Complete CPSs represent the coherent conditional beliefs a player would form if she asked herself this question for every possible event. This representation is used, e.g., in Myerson (1986) and Battigalli (1996).

In general, a richer class of conditioning events gives more bite to the chain rule, because each event is related by set inclusion to more events. Specifically, consider $\mathcal{C} \subseteq \mathcal{D}$ and let

$$\text{proj}_{\Delta^{\mathcal{C}}(\Omega)} \Delta^{\mathcal{D}}(\Omega) = \{\mu \in \Delta^{\mathcal{C}}(\Omega) : \exists \bar{\mu} \in \Delta^{\mathcal{D}}(\Omega), \forall C \in \mathcal{C}, \mu(\cdot|C) = \bar{\mu}(\cdot|C)\}$$

denote the set of CPSs on (Ω, \mathcal{C}) that can be derived from some CPS on (Ω, \mathcal{D}) .

Remark 2. If $\mathcal{C} \subseteq \mathcal{D}$, then $\text{proj}_{\Delta^{\mathcal{C}}(\Omega)} \Delta^{\mathcal{D}}(\Omega) \subseteq \Delta^{\mathcal{C}}(\Omega)$.

Thus, complete CPSs deserve special attention in that they embody the most stringent

restrictions on how players form conditional beliefs. Next we state a result providing a sufficient condition on collection \mathcal{C} so that $\text{proj}_{\Delta^{\mathcal{C}}(\Omega)}\Delta^*(\Omega) = \Delta^{\mathcal{C}}(\Omega)$, i.e., every CPS $\mu \in \Delta^{\mathcal{C}}(\Omega)$ has a **complete extension** $\bar{\mu} \in \Delta^*(\Omega)$ such that $\mu(\cdot|C) = \bar{\mu}(\cdot|C)$ for every $C \in \mathcal{C}$.

Definition 3. A collection $\mathcal{C} \subseteq 2^\Omega \setminus \{\emptyset\}$ is an **event tree** if (i) $\Omega \in \mathcal{C}$, and (ii) for all $C, D \in \mathcal{C}$, either $C \cap D = \emptyset$, or $C \subseteq D$, or $D \subseteq C$.

In words, an event tree represents possible paths of accumulation of information about the state of the world. For the analysis of strategic thinking it is important to consider systems of beliefs that assign probability 1 to some key events (such as the rationality of the co-players) *whenever possible* (cf. Battigalli 1996, Battigalli & Siniscalchi 2002). Say that $\mu \in \Delta^{\mathcal{C}}(\Omega)$ **strongly believes** event $E \neq \emptyset$ if, for every $C \in \mathcal{C}$,

$$E \cap C \neq \emptyset \Rightarrow \mu(E|C) = 1.$$

Lemma 1. Fix nonempty nested events $E^n \subseteq \dots \subseteq E^1$ and an event tree \mathcal{C} . For every CPS $\mu \in \Delta^{\mathcal{C}}(\Omega)$ that strongly believes each event E^1, \dots, E^n , there is a complete CPS $\bar{\mu} \in \Delta^*(\Omega)$ that extends μ and strongly believes each event E^1, \dots, E^n . In particular, $\text{proj}_{\Delta^{\mathcal{C}}(\Omega)}\Delta^*(\Omega) = \Delta^{\mathcal{C}}(\Omega)$.

3.2 Systems of beliefs

We now turn to the main object of our analysis: belief systems. We assume that player i would hold at each of her information sets $h_i \in H_i$ a belief $\mu^i(\cdot|h_i)$ assigning positive probability only to co-players' strategies that are compatible with h_i . Thus, players are endowed with **systems of beliefs** (often abbreviated in “**belief systems**”)

$$\mu^i = (\mu^i(\cdot|h_i))_{h_i \in H_i} \in \times_{h_i \in H_i} \Delta(S_{-i}(h_i)).$$

We discipline belief change with coherence properties. The first property is embodied in the definition of belief system: *knowledge implies probability-1 belief*, that is, players assign positive probability only to behavior of others consistent with their information. The second property concerns how players update or revise their beliefs across information sets; depending on these assumptions, we obtain different notions of belief systems. We consider three increasingly restrictive notions of consistency in updating/revision. The weakest one requires beliefs to be linked by the chain rule of conditional probabilities *as one moves forward* on a path, that is, if h'_i follows h_i and i deems h'_i reachable from h_i , then beliefs at h'_i should be derived from beliefs at h_i by conditioning on $S_{-i}(h'_i)$.

Definition 4. A system of beliefs μ^i is **forward consistent**, written $\mu^i \in \Delta_{\mathbb{F}^i}^{H_i}(S_{-i})$, if for all $h_i, h'_i \in H_i$ with $h_i \prec_i h'_i$, for all $E_{-i} \subseteq S_{-i}(h'_i)$

$$\mu^i(S_{-i}(h'_i)|h_i) > 0 \Rightarrow \mu^i(E_{-i}|h'_i) = \frac{\mu^i(E_{-i}|h_i)}{\mu^i(S_{-i}(h'_i)|h_i)}.$$

With a slight abuse of language, we often refer to the implication above as the “**forward chain rule**.” As discussed in the heuristic example of the Introduction, forward consistency allows players to form different beliefs at two information sets representing the same information on co-players’ behavior. This may happen if h_i and h'_i were both unexpected and differ only because of i ’s own behavior at the previous information set, so that $S_{-i}(h_i) = S_{-i}(h'_i)$. In this case, h_i and h'_i do not precede each other, and so forward consistency allows a player to revise beliefs differently at h_i and h'_i . This need not be viewed as a form of “irrationality.” A belief system can be interpreted as an external observer’s description of the beliefs a player would hold in every circumstance. Players may not be fully introspective, and know only their current beliefs. We clarify in Section 6 that to form rational plans, players need not plan in advance how they would behave at unexpected contingencies, and hence need not think in advance about what they would believe in such occurrences. Thus, players might revise their beliefs in different ways depending on their past actions.

Yet, one may still want to assume that players do not form different beliefs about others depending on their own behavior. Our intermediate restriction on belief systems excludes such cases by deriving belief systems from CPSs on (S_{-i}, \mathcal{H}_i) . CPSs embody by construction the property that same knowledge on others’ behavior implies same beliefs. Indeed, if $S_{-i}(h_i) = S_{-i}(h'_i)$, then h_i and h'_i correspond to the same element of \mathcal{H}_i , and hence for a CPS $\bar{\mu}^i$ we have by construction $\bar{\mu}^i(\cdot|S_{-i}(h_i)) = \bar{\mu}^i(\cdot|S_{-i}(h'_i))$.

Definition 5. A system of beliefs μ^i is **standard**, written $\mu^i \in \Delta^{H_i}(S_{-i})$, if there exists a CPS $\bar{\mu}^i \in \Delta^{\mathcal{H}_i}(S_{-i})$ such that $\mu^i(\cdot|h_i) = \bar{\mu}^i(\cdot|S_{-i}(h_i))$ for every $h_i \in H_i$.

(We call them “standard” because CPSs on (S_{-i}, \mathcal{H}_i) have been widely used in the literature on strategic reasoning since Battigalli & Siniscalchi 2002.)

Remark 3. A belief system μ^i is standard if and only if for all $h_i, h'_i \in H_i$ with $S_{-i}(h'_i) \subseteq S_{-i}(h_i)$, and for all $E_{-i} \subseteq S_{-i}(h'_i)$, we have

$$\mu^i(S_{-i}(h'_i)|h_i) > 0 \Rightarrow \mu^i(E_{-i}|h'_i) = \frac{\mu^i(E_{-i}|h_i)}{\mu^i(S_{-i}(h'_i)|h_i)}.$$

By perfect recall, $h_i \prec_i h'_i$ implies $S_{-i}(h'_i) \subseteq S_{-i}(h_i)$; thus, all standard belief systems are

forward consistent.

As already noted, H_i typically has a larger cardinality than \mathcal{H}_i because information sets may also represent information about player i 's behavior, not just the behavior of the co-players $-i$. For example, in game Γ of Figure 1, $H_i = \{\{\emptyset\}, \{(L, C)\}, \{(R, C)\}\}$ has three elements, while $\mathcal{H}_i = \{S_j, \{C.l, C.r\}\}$ has two elements. Despite this, the set of standard belief systems $\Delta^{H_i}(S_{-i})$ is isomorphic to $\Delta^{\mathcal{H}_i}(S_{-i})$. Indeed, for every standard belief system μ^i , $S_{-i}(h_i'') = S_{-i}(h_i')$ implies $\mu^i(\cdot|h_i'') = \mu^i(\cdot|h_i')$; with this, given a standard belief system μ^i , the array $(\bar{\mu}^i(\cdot|C))_{C \in \mathcal{H}_i}$ where $\bar{\mu}^i(\cdot|S_{-i}(h_i)) = \mu^i(\cdot|h_i)$ for each $h_i \in H_i$ defines uniquely a CPS on (S_{-i}, \mathcal{H}_i) . So $\mu^i(\cdot|h_i) \leftrightarrow \bar{\mu}^i(\cdot|S_{-i}(h_i))$ is a bijection between $\Delta^{H_i}(S_{-i})$ and $\Delta^{\mathcal{H}_i}(S_{-i})$. In Section 4, we characterize standard belief systems in games with observable actions.

The strongest notion of consistency requires a belief system to be induced by some complete CPS.

Definition 6. A system of beliefs μ^i is **completely consistent**, written $\mu^i \in \Delta_C^{H_i}(S_{-i})$, if there exists a complete CPS $\bar{\mu}^i \in \Delta^*(S_{-i})$ such that $\mu^i(\cdot|h_i) = \bar{\mu}^i(\cdot|S_{-i}(h_i))$ for every $h_i \in H_i$.

By Remark 2, complete consistency may be more restrictive than standard consistency: $\Delta_C^{H_i}(S_{-i}) \subseteq \Delta^{H_i}(S_{-i})$. Yet, the two notions coincide if the collection of conditioning events \mathcal{H}_i is an event tree:

Remark 4. If \mathcal{H}_i is an event tree, then $\Delta^{H_i}(S_{-i}) = \Delta_C^{H_i}(S_{-i})$.

Proof: By Lemma 1, $\Delta^{\mathcal{H}_i}(S_{-i}) = \text{proj}_{\Delta^{\mathcal{H}_i}(S_{-i})} \Delta^*(S_{-i})$, which implies $\Delta^{H_i}(S_{-i}) = \Delta_C^{H_i}(S_{-i})$. ■

If, instead, \mathcal{H}_i is not an event tree, then Lemma 1 does not apply and the inclusion may be strict. The following ‘‘common ratio’’ property gives a sense of the strength of complete consistency.

Remark 5. If belief system μ^i is completely consistent, then, for all $h_i, h_i' \in H_i$ and $s_{-i}, t_{-i} \in S_{-i}(h_i) \cap S_{-i}(h_i')$,

$$\mu^i(t_{-i}|h_i), \mu^i(t_{-i}|h_i') > 0 \Rightarrow \frac{\mu^i(s_{-i}|h_i)}{\mu^i(t_{-i}|h_i)} = \frac{\mu^i(s_{-i}|h_i')}{\mu^i(t_{-i}|h_i')}.$$

Proof: Consider a belief system μ^i derived from some complete CPS $\bar{\mu}^i$. Fix any $h_i, h_i' \in H_i$ and $s_{-i}, t_{-i} \in S_{-i}(h_i) \cap S_{-i}(h_i')$ such that $\mu^i(t_{-i}|h_i) > 0$ and $\mu^i(t_{-i}|h_i') > 0$. Since $\bar{\mu}^i$ is a

complete CPS, the chain rule relates $\bar{\mu}^i(\cdot|S_{-i}(h_i))$ and $\bar{\mu}^i(\cdot|S_{-i}(h'_i))$ to $\bar{\mu}^i(\cdot|\{s_{-i}, t_{-i}\})$. Hence,

$$\mu^i(s_{-i}|h_i) = \bar{\mu}^i(s_{-i}|S_{-i}(h_i)) = \bar{\mu}^i(s_{-i}|\{s_{-i}, t_{-i}\}) \cdot \bar{\mu}^i(\{s_{-i}, t_{-i}\}|S_{-i}(h_i)),$$

$$\mu^i(s_{-i}|h'_i) = \bar{\mu}^i(s_{-i}|S_{-i}(h'_i)) = \bar{\mu}^i(s_{-i}|\{s_{-i}, t_{-i}\}) \cdot \bar{\mu}^i(\{s_{-i}, t_{-i}\}|S_{-i}(h'_i)).$$

Since $\mu^i(t_{-i}|h_i), \mu^i(t_{-i}|h'_i) > 0$ we have $\bar{\mu}^i(\{s_{-i}, t_{-i}\}|S_{-i}(h_i)) = \mu^i(\{s_{-i}, t_{-i}\}|h_i) > 0$ and so

$$\frac{\mu^i(s_{-i}|h_i)}{\mu^i(\{s_{-i}, t_{-i}\}|h_i)} = \bar{\mu}^i(s_{-i}|\{s_{-i}, t_{-i}\}) = \frac{\mu^i(s_{-i}|h'_i)}{\mu^i(\{s_{-i}, t_{-i}\}|h'_i)}.$$

It can be analogously verified for t_{-i} that

$$\frac{\mu^i(t_{-i}|h_i)}{\mu^i(\{s_{-i}, t_{-i}\}|h_i)} = \bar{\mu}^i(t_{-i}|\{s_{-i}, t_{-i}\}) = \frac{\mu^i(t_{-i}|h'_i)}{\mu^i(\{s_{-i}, t_{-i}\}|h'_i)}.$$

The last two equations yield, dividing the former by the latter,

$$\frac{\mu^i(s_{-i}|h_i)}{\mu^i(t_{-i}|h_i)} = \frac{\mu^i(s_{-i}|h'_i)}{\mu^i(t_{-i}|h'_i)}.$$

■

Standard belief systems need not satisfy this property, as shown in Example 1 of Section 4. Yet, this necessary condition, although quite strong, is not sufficient for complete consistency. We provide in Section 5 an example of this fact, as well as the exact characterization of complete consistency in terms of coherence of odds ratios at different information sets. It is also worth noting that complete consistency is connected to Kreps & Wilson's (1982) consistency in the following sense:

Remark 6. A system of beliefs μ^i is completely consistent if and only if there is a sequence of strictly positive probability measures $(\nu_n)_{n \in \mathbb{N}}$ on S_{-i} such that

$$\forall h_i \in H_i, \forall s_{-i} \in S_{-i}(h_i), \mu^i(s_{-i}|h_i) = \lim_{n \rightarrow \infty} \frac{\nu_n(s_{-i})}{\nu_n(S_{-i}(h_i))}.$$

This is implied by Remark 12 in the Appendix, which draws on standard results about complete CPSs (see, e.g., Myerson 1986). Kreps & Wilson (1982) put forward an across-players notion of consistency of assessments derived from Selten's (1975) "trembling hand" idea. Remark 6 implies that, in two-person games, complete consistency is equivalent to a single-player version of Kreps & Wilson's consistency. See the Appendix for details.

To summarize, we defined three nested subsets of belief systems (the inclusions may be

strict):

$$\Delta_{\mathbf{C}}^{H_i}(S_{-i}) \subseteq \Delta^{H_i}(S_{-i}) \subseteq \Delta_{\mathbf{F}}^{H_i}(S_{-i}).$$

In the next sections, we shed further light on the interpretation of these belief systems by providing a characterization of standard consistency in games with observable actions and a characterization of complete consistency. Then we move on to rational planning and strategic reasoning as captured by rationalizability and we show that, in spite of a considerable difference in the degree of consistency assumed by these types of belief systems, these differences are inconsequential for rationalizability.

4 Standard belief systems in games with observable actions

In this section, we provide a characterization of standard belief systems in *games with observable actions*. We show that a belief system is standard if and only if it is forward consistent and it assigns the same beliefs at histories that differ only in player i 's own action at the immediate predecessor, i.e., beliefs about co-players' strategies are independent of own behavior. We then argue that standard and forward consistent belief systems coincide in games with perfect information. Understanding the relationship between forward consistent and standard belief systems is relevant because, as anticipated in the Introduction and shown in Sections 8.4-8.5, it may matter for predictions that rely on strategic reasoning under contextual restrictions on conditional beliefs.

It is first convenient to study the structure of the “strategic-form information sets” $S_{-i}(h)$ ($h \in H$) in games with observable actions. In such games, non-terminal histories and information sets essentially coincide; therefore, for every $h \in H$, we have $S(h) = \times_{i \in I} S_i(h)$. For each $i \in I$, $h \in H$, and $a_{-i} \in \mathcal{A}_{-i}(h)$, let

$$S_{-i}(h, a_{-i}) = \{s_{-i} \in S_{-i}(h) : s_{-i}(h) = a_{-i}\}$$

denote the set of co-players' strategy profiles consistent with h that select a_{-i} at h . For any $h', h'' \in H$, let $\pi(h', h'') \in H$ be the **last common predecessor** (longest common prefix) of h' and h'' . For any $h \prec h'$, let $\alpha(h, h') = (\alpha_j(h, h'))_{j \in I} \in \mathcal{A}(h)$ be the unique action profile such that $(h, \alpha(h, h')) \preceq h'$, and let $f(h, h')$ be the **immediate follower** of h (weakly) **preceding** h' . Note, by definition, $S_{-i}(f(h, h')) = S_{-i}(h, \alpha_{-i}(h, h'))$. The following lemma is crucial to characterize standard belief systems.

Lemma 2. *In games with observable actions, if two histories $h', h'' \in H$ are not ordered by precedence and \bar{h} is their longest common predecessor, then $S_{-i}(h') \subseteq S_{-i}(h'')$ implies that*

(a) $\alpha_{-i}(\bar{h}, h') = \alpha_{-i}(\bar{h}, h'')$ and (b) $S_{-i}(h'') = S_{-i}(f(\bar{h}, h'')) = S_{-i}(f(\bar{h}, h'))$.

Proof: Take any two histories $h', h'' \in H$ not ordered by precedence and consider their last common predecessor $\bar{h} = \pi(h', h'')$. We first prove by contraposition that, if $S_{-i}(h') \subseteq S_{-i}(h'')$, then only player i can be active at histories h such that $\bar{h} \prec h \prec h''$. Indeed, suppose that some player $j \neq i$ is active at such h ; then, j has an action $a_j^* \in \mathcal{A}_j(h) \setminus \{\alpha_j(h, h'')\}$, which implies that there exists $a_{-i}^* \in \mathcal{A}_{-i}(h)$ such that $a_{-i}^* \neq \alpha_{-i}(h, h'')$. Take any $s_{-i}^* \in S_{-i}$ such that (i) for all $\tilde{h} \prec h'$, $s_{-i}(\tilde{h}) = \alpha_{-i}(\tilde{h}, h')$, so that $s_{-i}^* \in S_{-i}(h')$, and (ii) $s_{-i}(h) = a_{-i}^*$, so that $s_{-i}^* \notin S_{-i}(h'')$ (noting that $h \not\prec h''$, this does not conflict with (i)). With this, there exists $s_{-i}^* \in S_{-i}(h') \setminus S_{-i}(h'')$, that is, $S_{-i}(h') \not\subseteq S_{-i}(h'')$.

Now suppose that $S_{-i}(h') \subseteq S_{-i}(h'')$. Then $\alpha_{-i}(\bar{h}, h') = \alpha_{-i}(\bar{h}, h'')$, otherwise $S_{-i}(h')$ and $S_{-i}(h'')$ would be disjoint. This implies $S_{-i}(f(\bar{h}, h')) = S_{-i}(f(\bar{h}, h''))$. Since only player i can be active at histories h with $\bar{h} \prec h \prec h''$, we have $S_{-i}(h'') = S_{-i}(f(\bar{h}, h''))$. ■

Using this result, we can characterize standard belief systems by means of a property of independence of beliefs from own behavior.

Definition 7. A system of beliefs μ^i satisfies **own-action independence (OI)** if $\mu^i(\cdot|h') = \mu^i(\cdot|h'')$ for all $h' = (h, a')$ and $h'' = (h, a'')$ with $a'_{-i} = a''_{-i}$.

Denote by $\Delta_{\text{F,OI}}^H(S_{-i})$ and $\Delta_{\text{OI}}^H(S_{-i})$ the sets of forward consistent and standard belief systems that satisfy own-action independence.

Remark 7. In games with observable actions, all standard belief systems satisfy own-action independence, i.e., $\Delta_{\text{OI}}^H(S_{-i}) = \Delta^H(S_{-i})$. Indeed, if $h' = (h, a')$ and $h'' = (h, a'')$ are such that $a'_{-i} = a''_{-i}$, then $S_{-i}(h') = S_{-i}(h'')$. Then by Remark 3 the chain rule applies to the pair (h', h'') , which implies, for equal conditioning events, equal conditional measures.

Theorem 1. *In games with observable actions, a system of beliefs is standard if and only if it is forward consistent and satisfies OI, that is, $\Delta^H(S_{-i}) = \Delta_{\text{F,OI}}^H(S_{-i})$.*

Proof: The inclusion $\Delta^H(S_{-i}) = \Delta_{\text{OI}}^H(S_{-i}) \subseteq \Delta_{\text{F,OI}}^H(S_{-i})$ is obvious. We show the other inclusion. Take any $\mu^i \in \Delta_{\text{F,OI}}^H(S_{-i})$ and two histories $h', h'' \in H$ with $S_{-i}(h') \subseteq S_{-i}(h'')$. By Remark 3, it is enough to show that the forward chain rule relates $\mu^i(\cdot|h'')$ to $\mu^i(\cdot|h')$. If $h'' \preceq h'$ the forward chain rule relates $\mu^i(\cdot|h'')$ to $\mu^i(\cdot|h')$. If h'' and h' are not related by precedence, then let $h = \pi(h', h'')$, $a' = \alpha(h, h')$ and $a'' = \alpha(h, h'')$. Since $S_{-i}(h') \subseteq S_{-i}(h'')$, Lemma 2 implies (a) $a'_{-i} = a''_{-i}$ and (b) $S_{-i}(h'') = S_{-i}((h, a'')) = S_{-i}((h, a'))$. Since $(h, a'') \preceq h''$, the forward chain rule relates $\mu^i(\cdot|(h, a''))$ to $\mu^i(\cdot|h'')$. Since $S_{-i}(h'') = S_{-i}((h, a''))$, this implies $\mu^i(\cdot|h'') = \mu^i(\cdot|(h, a''))$. Since $a'_{-i} = a''_{-i}$ and μ^i satisfies OI, we have $\mu^i(\cdot|(h, a'')) = \mu^i(\cdot|(h, a'))$. Finally, since $(h, a') \preceq h'$, the forward chain rule relates $\mu^i(\cdot|(h, a'))$ to $\mu^i(\cdot|h')$.

Thus, $\mu^i(\cdot|h'') = \mu^i(\cdot|(h, a'')) = \mu^i(\cdot|(h, a'))$ and the forward chain rule indirectly relates $\mu^i(\cdot|h'')$ to $\mu^i(\cdot|h')$. ■

This characterization is useful in constructive proofs when working with standard beliefs systems in games with observable actions: it is sometimes easier to show that a system of beliefs satisfies the forward chain rule and own-action independence than showing directly standard consistency.

On one hand, own-action independence of beliefs is key to be able to conclude that $\mu^i(\cdot|(h, a')) = \mu^i(\cdot|(h, a''))$ in the proof. Otherwise, if $\mu^i(S_{-i}(h, a'_{-i})|h) = 0$, forward consistency does not guarantee that i revises her beliefs at (h, a') and (h, a'') in the same way and hence we would not be guaranteed that the forward chain rule indirectly relate $\mu^i(\cdot|h')$ and $\mu^i(\cdot|h'')$.

On the other hand, a forward consistent belief system μ^i may violate own-action independence of beliefs at two histories $h' = (h, a')$, $h'' = (h, a'')$ with $a'_{-i} = a''_{-i}$ only if (a) $a'_i \neq a''_i$, which implies that i is active at h , and (b) $\mu^i(S_{-i}(h, a'_{-i})|h) = 0$, which implies that at least one co-player of i is active at h , since otherwise $S_{-i}(h, a'_{-i}) = S_{-i}(h)$ and thus $\mu^i(S_{-i}(h, a'_{-i})|h) = \mu^i(S_{-i}(h)|h) = 1 > 0$. So, forward consistent belief systems might feature own-action *dependence* only in games where there are simultaneous moves at some stage, which yields the following result:¹¹

Corollary 1. *In games with perfect information, a system of beliefs is standard if and only if it is forward consistent: $\Delta^H(S_{-i}) = \Delta_F^H(S_{-i})$.*

Since the foregoing results are independent of players' interactive knowledge of the profile of payoff functions, they extend to games with incomplete information and observable actions.

As for complete consistency, it is typically stronger than forward/standard consistency, even in games with perfect information, as shown by the example below.

Example 1. Consider histories $h' = (L, d')$ and $h'' = (R, d'')$ in the perfect-information game structure depicted in Figure 4. For strategies $s_{-i} = d'.d''.\ell$, $s'_{-i} = d'.d''.r$, and $s^*_{-i} = a'.a''.\ell$ (where $-i = j$ is Joe), we have $s_{-i}, s'_{-i} \in S_{-i}(h') \cap S_{-i}(h'')$, $s^*_{-i} \notin S_{-i}(h')$, and $s^*_{-i} \notin S_{-i}(h'')$. This shows that \mathcal{H}_i is not an event tree and Remark 4 does not apply. Take a forward consistent (hence, by Corollary 1, *standard*) belief system μ^i of Isa such that $\mu^i(s^*_{-i}|\emptyset) = 1$. Isa's beliefs after L and R coincide with her initial belief, because Joe is not active at the root. It follows that $\mu^i(S_{-i}(h')|L) = 0$ and $\mu^i(S_{-i}(h'')|R) = 0$, so that Isa has to revise her beliefs at both h' and h'' . Since $S_{-i}(h')$ and $S_{-i}(h'')$ are different and not nested, we can set $\mu^i(s_{-i}|h') = 1$ and $\mu^i(s'_{-i}|h'') = 1$. But this violates the odds-ratio property of

¹¹Ben-Porath (1997) analyzes perfect-information games using forward consistent belief systems. Corollary 1 implies that, implicitly, he is using standard belief systems.




Figure 4: A perfect-information game tree.

Remark 5, which would require in this specific case that $\mu^i(s_{-i}|h') = \mu^i(s_{-i}|h'')$. Thus, μ^i is a standard belief system that is not completely consistent. \blacktriangle

5 A characterization of complete consistency

In this section, we show that complete consistency is characterized by a strong coherence condition relating probability ratios across different information sets.

We already noted that complete consistency implies constant odds ratios between pair of opponents' strategy profiles at all information sets consistent with both strategy profiles. Yet, this condition is not sufficient for complete consistency, as shown by the following example.

Example 2. Consider the game structure depicted in Figure 5. Note that there is no pair (h_i, h'_i) of third-stage information sets of Isa such that $S_j(h_i) \cap S_j(h'_i)$ contains two distinct actions/strategies of Joe. Thus, if belief system μ^i is forward consistent it also satisfies the “odds-ratio” property of Remark 5, which applies only if either h_i or h'_i , is the initial information set of Isa. Suppose that Isa is initially certain of termination, that is, $\mu^i(t|\emptyset) = 1$, then forward consistency and the “odds-ratios” property are trivially satisfied by μ^i . Yet, μ^i need not be completely consistent. To see this, let $[xy] = \{(xy, x), (xy, y)\}$ denote the third-stage information set of Isa after she chose xy and Joe chose either x or y (both different from t). Since $\mu^i(t|\emptyset) = 1$, forward consistency puts no constraints on beliefs at such information sets. Then we may have $\mu^i(a|[ab]) > \frac{1}{2}$, $\mu^i(b|[bc]) > \frac{1}{2}$, and $\mu^i(c|[ac]) > \frac{1}{2}$, which violates complete consistency. Indeed, any complete CPS $\bar{\mu}^i \in \Delta^*(S_j)$ defines i 's belief conditional on $\{a, b, c\} \subset S_j$, and hence the chain rule relates beliefs conditional on $\{a, b\}$, $\{b, c\}$ and $\{a, c\}$ to $\bar{\mu}^i(\cdot|\{a, b, c\})$. \blacktriangle




Figure 5: Complete consistency is stronger than the “odds-ratio property”.

A standard odds ratio is computed using the conditional probabilities of two elements s_{-i} and s'_{-i} at the same information set. As such, it does not allow to compare the probabilities of s_{-i} and s'_{-i} at different information sets. This comparison is made possible by the following notion of “generalized odds ratio,” introduced by Catonini (2022).

Definition 8. [Catonini 2022] Given a belief system $(\mu^i(\cdot|h_i))_{h_i \in H_i}$ and a pair $s_{-i}, s'_{-i} \in S_{-i}$, a (possibly infinite) **generalized odds ratio** between s_{-i} and s'_{-i} is the product of a finite concatenation of odds ratios, namely:

$$q = \frac{\mu^i(s_{-i}|h_i^1)}{\mu^i(s'_{-i}|h_i^1)} \cdot \frac{\mu^i(s_{-i}^1|h_i^2)}{\mu^i(s'_{-i}^1|h_i^2)} \cdot \dots \cdot \frac{\mu^i(s_{-i}^{n-1}|h_i^n)}{\mu^i(s'_{-i}^{n-1}|h_i^n)},$$

where, for all $k \in \{1, \dots, n\}$, (i) $s_{-i}^{k-1}, s_{-i}^k \in S_{-i}(h_i^k)$ (with $s_{-i}^0 = s_{-i}$ and $s_{-i}^n = s'_{-i}$), and (ii) if $\mu^i(s_{-i}^k|h_i^k) = 0$, then $\mu^i(s_{-i}^{k'-1}|h_i^{k'}) \neq 0$ for all $k' \in \{1, \dots, n\}$.

For instance, if player i deems s_{-i} twice more likely than another strategy s'_{-i} at some information set where both s_{-i} and s'_{-i} are still possible, and deems s'_{-i} twice more likely than s''_{-i} at some other information set where s'_{-i} and s''_{-i} are still possible, then we can say that there is a generalized odds ratio of 4 between s_{-i} and s''_{-i} in player i 's belief system.

A generalized odds ratio can be zero or infinite (but not indeterminate). A zero generalized odds ratio between s_{-i} and s'_{-i} can be given the following interpretation: If player i was to make an indirect assessment of the relative likelihood of s_{-i} and s'_{-i} by looking at the relative likelihoods between these and other strategies at different information sets, she must conclude

that s_{-i} has probability zero whenever s'_{-i} is possible. Whether she actually does so depends on the degree of consistency in her belief system. Indeed, Catonini (2022) shows that complete consistency is equivalent to the condition that, for any pair of co-players' strategy profiles, all generalized odds ratios are identical.¹²

Theorem 2. (Catonini 2022) *A belief system μ^i is completely consistent if and only if for every pair $s_{-i}, s'_{-i} \in S_{-i}$, all generalized odds ratios of μ^i between s_{-i} and s'_{-i} are identical.*

Here we provide a direct proof of the fact that complete consistency requires all generalized odds ratios between any two strategies to coincide.

Proof of the “only if” part. We provide a direct proof that complete consistency implies constant GORs. Fix a pair (s_{-i}, s'_{-i}) and consider two GORs

$$\begin{aligned} q^1 &= \frac{\mu^i(s_{-i}|h_i^1)}{\mu^i(s'_{-i}|h_i^1)} \cdot \frac{\mu^i(s_{-i}^1|h_i^2)}{\mu^i(s'_{-i}|h_i^2)} \cdot \dots \cdot \frac{\mu^i(s_{-i}^{n-1}|h_i^n)}{\mu^i(s'_{-i}|h_i^n)}, \\ q^2 &= \frac{\mu^i(s_{-i}|h_i^m)}{\mu^i(s_{-i}^{m-1}|h_i^m)} \cdot \frac{\mu^i(s_{-i}^{m-1}|h_i^{m-1})}{\mu^i(s_{-i}^{m-2}|h_i^{m-1})} \cdot \dots \cdot \frac{\mu^i(s_{-i}^{n+1}|h_i^{n+1})}{\mu^i(s'_{-i}|h_i^{n+1})}, \end{aligned}$$

where we use a labeling convention so that $m > n$. We want to show that $q^1 = q^2$. If they are both zero or both infinite, we are done. Otherwise, suppose without loss of generality that $q^2 \geq q^1$, so that $q^1 \neq \infty$ and $q^2 > 0$. Calling $s_{-i}^n = s'_{-i}$ and $s_{-i}^0 = s_{-i}$, let

$$q := \frac{q^1}{q^2} = \frac{\mu^i(s_{-i}^0|h_i^1)}{\mu^i(s_{-i}^1|h_i^1)} \cdot \dots \cdot \frac{\mu^i(s_{-i}^{n-1}|h_i^n)}{\mu^i(s_{-i}^n|h_i^n)} \cdot \frac{\mu^i(s_{-i}^n|h_i^{n+1})}{\mu^i(s_{-i}^{n+1}|h_i^{n+1})} \cdot \dots \cdot \frac{\mu^i(s_{-i}^{m-1}|h_i^m)}{\mu^i(s_{-i}^0|h_i^m)}.$$

Thus, all denominators are positive. Clearly, $q^1 = q^2$ if $q = 1$. To show $q = 1$, we proceed as follows: Let $\bar{\mu}^i = (\bar{\mu}^i(\cdot|h_i))_{2^{S_{-i} \setminus \{\emptyset\}}}$ be a complete CPS that extends μ^i . Let

$$C = \cup_{k=1}^m S_{-i}(h_i^k).$$

Since $\bar{\mu}^i(C|C) = 1$, there is $k \in \{1, \dots, m\}$ such that $\bar{\mu}^i(S_{-i}(h_i^k)|C) > 0$. Since $\bar{\mu}^i(s_{-i}^k|S_{-i}(h_i^k)) = \mu^i(s_{-i}^k|h_i^k) > 0$, by the chain rule $\bar{\mu}^i(s_{-i}^k|C) > 0$. By definition of the odds ratios, $s_{-i}^k \in S_{-i}(h_i^{k+1})$ (whenever $k+1 \leq m$), hence $\bar{\mu}^i(S_{-i}(h_i^{k+1})|C) > 0$. Expanding this reasoning we obtain that $\bar{\mu}^i(S_{-i}(h_i^p)|C) > 0$ for every $p = k, \dots, m$. Now, note that $s_{-i}^m = s_{-i}^0$, and that $s_{-i}^0 \in S_{-i}(h_i^1)$. So, $\bar{\mu}^i(S_{-i}(h_i^m)|C) > 0$ implies $\bar{\mu}^i(s_{-i}^0|C) > 0$ by the chain rule with $\bar{\mu}^i(\cdot|S_{-i}(h_i^m))$ and $\bar{\mu}^i(s_{-i}^0|S_{-i}(h_i^m)) > 0$. Thus, $\bar{\mu}^i(S_{-i}(h_i^1)|C) > 0$, and repeating the above

¹²The (2020) working-paper version of Siniscalchi (2022) provides an alternate characterization of complete consistency that is closely related to Catonini's. The connection is formalized in Catonini (2022).

argument we have $\bar{\mu}^i(S_{-i}(h_i^p)|C) > 0$ for all $p = 1, \dots, m$. Therefore, by the chain rule

$$\frac{\mu^i(s_{-i}^{p-1}|h_i^p)}{\mu^i(s_{-i}^p|h_i^p)} = \frac{\bar{\mu}^i(s_{-i}^{p-1}|C)}{\bar{\mu}^i(s_{-i}^p|C)}.$$

Hence,

$$q = \frac{\bar{\mu}^i(s_{-i}^0|C)}{\bar{\mu}^i(s_{-i}^1|C)} \cdot \frac{\bar{\mu}^i(s_{-i}^1|C)}{\bar{\mu}^i(s_{-i}^2|C)} \cdot \dots \cdot \frac{\bar{\mu}^i(s_{-i}^{m-1}|C)}{\bar{\mu}^i(s_{-i}^0|C)} = 1.$$

■

Thus, complete consistency assumes a very high degree of coherence concerning how players revise their beliefs at different information sets. Yet this coherence need not be viewed as a consequence of conscious calculations of a fully introspective player about his conditional beliefs. A player may still be partially introspective about how she would revise her beliefs at unexpected contingencies, yet complete consistency could be a “wired-in” property of belief formation. Catonini (2022) relates complete consistency to the possibility of being “Dutch-booked in objective expected terms,” suggesting that individuals violating complete consistency may be prone to undertaking detrimental courses of actions in the long run.

6 Rational planning

In this section, we define standard criteria of optimality for strategies given a system of beliefs and develop an interpretation in terms of “partial planning,” which does not require players to be fully introspective about their conditional beliefs, but only requires players to know their current beliefs. We then relate this interpretation to the degree of consistency that belief systems may feature.

Fix a strategy s_i and a belief system μ^i . Recall that $H_i(s_i)$ is the collection of information sets that may occur if i implements strategy s_i . For any $h_i \in H_i(s_i)$, the conditional expected payoff of s_i given h_i under μ^i is

$$U_i(s_i, \mu^i(\cdot|h_i)) = \sum_{s'_{-i} \in S_{-i}(h_i)} u_i(\zeta(s_i, s'_{-i})) \mu^i(s'_{-i}|h_i).$$

Remark 8. Fix s_i , μ^i , and $h_i \in H_i(s_i)$ arbitrarily. By Remark 1, expected payoff $U_i(s_i, \mu^i(\cdot|h_i))$ is independent of how s_i is specified at information sets that cannot occur under s_i : $U_i(s_i, \mu^i(\cdot|h_i)) = U_i(s'_i, \mu^i(\cdot|h_i))$ for every s'_i that selects the same actions as s_i on the sub-domain $H_i(s_i)$, that is, for every $s'_i \equiv_i s_i$.

For every $h_i \in H_i$, let $\rho_i[s_i/h_i]$ denote the “ h_i -replacement” strategy that, at every

$h'_i \prec_i h_i$, chooses the unique action $\alpha_i(h'_i, h_i)$ leading from h'_i toward h_i , and coincides with s_i at all other information sets, that is,

$$\forall h'_i \in H_i, \rho_i[s_i/h_i](h'_i) = \begin{cases} \alpha_i(h'_i, h_i) & \text{if } h'_i \prec_i h_i, \\ s_i(h'_i) & \text{if } h'_i \not\prec_i h_i. \end{cases}$$

By definition, $\rho_i[s_i/h_i] \in S_i(h_i)$ and $U_i(\rho_i[s_i/h_i], \mu^i(\cdot|h_i))$ is well posed for all $s_i \in S_i$ and $h_i \in H_i$. The most demanding notion of optimality of behavior consists in planning and implementing a strategy that maximizes expected payoff starting from every information set.

Definition 9. A strategy \bar{s}_i is **sequentially optimal** under belief system μ^i , written $\bar{s}_i \in BR_i^*(\mu^i)$, if

$$\forall h_i \in H_i, \forall s'_i \in S_i(h_i), \quad U_i(\rho_i[\bar{s}_i/h_i], \mu^i(\cdot|h_i)) \geq U_i(s'_i, \mu^i(\cdot|h_i)).$$

Note that, if belief system μ^i is not forward consistent, then the set of sequentially optimal strategies under μ^i may be empty because of conflicts between expected-payoff maximization at different ordered information sets.

Example 3. In game Γ''' of Figure 3, the chain rule implies that Isa must hold the same belief about Joe at the first and second node, because at both nodes she has no information about Joe. Thus, if she initially believes $\mu^i(\ell|\emptyset) > \frac{1}{2}$ and the chain rule holds, strategy $c'.c''$ is sequentially optimal. If instead $\mu^i(\ell|\emptyset) > \frac{1}{2}$ and $\mu^i(\ell|c') < \frac{1}{4}$, violating the chain rule, then the maximization problem at the root is still solved by $c'.c''$, but the one at history/node (c') is solved by $c'.d''$, and no strategy is sequentially optimal under μ^i . \blacktriangle

The aforementioned issue of conflicting conditional preferences is circumvented by an “intra-personal equilibrium” property that can always be satisfied in finite games with perfect recall: one-step optimality. For any strategy $s_i \in S_i$, information set $h_i \in H_i$, and action $a_i \in \mathcal{A}_i(h_i)$, let $\lambda_i[s_i/h_i, a_i]$ denote the “local (h_i, a_i) -replacement” strategy of i that selects a_i at h_i and coincides with $\rho_i[s_i/h_i]$ at all other information sets of i , that is,

$$\forall h'_i \in H_i, \lambda_i[s_i/h_i, a_i](h'_i) = \begin{cases} a_i & \text{if } h'_i = h_i, \\ \alpha_i(h'_i, h_i) & \text{if } h'_i \prec_i h_i, \\ s_i(h'_i) & \text{if } h'_i \not\prec_i h_i. \end{cases}$$

Definition 10. A strategy $\bar{s}_i \in S_i$ is **one-step optimal** under belief system μ^i if

$$\forall h_i \in H_i, \forall a_i \in \mathcal{A}_i(h_i), \quad U_i(\rho_i[\bar{s}_i/h_i], \mu^i(\cdot|h_i)) \geq U_i(\lambda_i[\bar{s}_i/h_i, a_i], \mu^i(\cdot|h_i)).$$

One-step optimality can be interpreted as the result of a **folding-back planning** algorithm: A fully introspective player i who knows her belief system μ^i first considers the information sets h_i with height 0 in H_i , that is, those where she makes a last move;¹³ for any such information set h_i , she plans to choose an expected-payoff-maximizing action given belief $\mu^i(\cdot|h_i)$ about the co-players (breaking ties arbitrarily). In step $\ell > 0$ of the algorithm, as she considers any information set h_i with height ℓ in H_i ,¹⁴ she plans to choose an expected-payoff-maximizing action given belief $\mu^i(\cdot|h_i)$ and the prediction that she would behave as already planned in earlier steps of the algorithm at the following information sets of height $k < \ell$ (again, breaking ties arbitrarily). Note that this algorithm is equivalent to one-step optimality even if the chain rule does not hold.

Example 4. Refer to game Γ''' . If $\mu^i(\ell|\emptyset) > \frac{1}{2}$ and $\mu^i(\ell|c') < \frac{1}{4}$, the only one-step optimal strategy is $d'.d''$: Isa understands how she would change her beliefs and that she would choose d'' at the second node, if reached; thus, she chooses d' at the root. In other words, d'' is just a conditional prediction of Isa about herself, not something that she initially deems optimal conditional on c' . \blacktriangle

Finiteness and perfect recall imply that a one-step optimal (pure) strategy can always be found by folding back:

Remark 9. For every belief system μ^i , at least one strategy is one-step optimal under μ^i .

When the chain rule holds, conditional preferences are dynamically consistent and—by a relatively standard dynamic programming argument—one can show that a strategy is sequentially optimal if and only if it can be obtained by folding-back planning. This implies that sequential and one-step optimality are equivalent, a result known as the *one-shot-deviation principle*. Notably, Perea (2002) proves that the one-shot-deviation principle holds also for the weakest form of chain rule considered here, forward consistency.¹⁵

Theorem 3. (Perea 2002) *For every forward consistent belief system μ^i , a strategy is sequentially optimal under μ^i if and only if it is one-step optimal under μ^i .*

Clearly, since $\Delta_C^{H_i}(S_{-i}) \subseteq \Delta^{H_i}(S_{-i}) \subseteq \Delta_F^{H_i}(S_{-i})$, the one-shot deviation principle also holds for standard and completely consistent belief systems. Intuitively, the reason why the

¹³For the sake of this discussion, the distinction between information sets where i is active or inactive is immaterial.

¹⁴One that is followed by at most ℓ moves of hers.

¹⁵Perea shows the result for two-person games, but the extension to n -person games is straightforward. The proof is available upon request.

equivalence between sequential and one-step optimality holds—despite possible violations of strong versions of the chain rule—is the following. Under a forward consistent belief system, player i may hold conflicting beliefs at information sets h'_i and h''_i *only if* they follow (a) different actions of hers and (b) unexpected actions of the co-players (consider $h'_i = \{(L, C)\}$ and $h''_i = \{(R, C)\}$ for Isa in game Γ' , if she initially assigns probability 0 to action C of Joe). In this case, the decision problems of i at h'_i and h''_i are mutually independent; furthermore, from the perspective of earlier information sets (the root for Isa in Γ'), planning at h'_i and h''_i does not affect expected payoffs, because i deems both h'_i and h''_i unreachable.

Example 5. Go back to Γ' of Figure 1. Let $\Delta_{a',a''}$ denote the set of belief systems μ^i such that $\mu^i(\{Q.l, Q.r\}|\emptyset) = 1$, $\mu^i(C.l|(L, C)) \leq \frac{2}{5}$, and $\mu^i(C.l|(R, C)) \geq \frac{1}{2}$. Strategy $R.a'.a''$ of Isa is one-step optimal under any $\mu^i \in \Delta_{a',a''}$: $\mu^i(\{Q.l, Q.r\}|\emptyset) = 1$ makes R a best reply at the root, $\mu^i(C.l|(L, C)) \leq \frac{2}{5}$ makes a' a best reply given (L, C) , and $\mu^i(C.l|(R, C)) \geq \frac{1}{2}$ makes a'' a best reply given (R, C) . Furthermore, all these belief systems are forward consistent, that is, $\Delta_{a',a''} \subseteq \Delta_{\mathbb{F}}^{H_i}(S_{-i})$, because $\mu^i(\{Q.l, Q.r\}|\emptyset) = 1$ implies that the forward chain rule holds trivially. Thus, Theorem 3 implies that $R.a'.a''$ is sequentially optimal under every $\mu^i \in \Delta_{a',a''}$. \blacktriangle

Theorem 3 and Remark 9 imply the following:

Corollary 2. *For every forward consistent belief system μ^i , at least one strategy is sequentially optimal under μ^i , that is, $BR_i^*(\mu^i) \neq \emptyset$.*

We now turn to a seemingly less demanding notion of optimality.

Definition 11. A strategy \bar{s}_i is **weakly sequentially optimal** under belief system μ^i , written $\bar{s}_i \in BR_i(\mu^i)$, if

$$\forall h_i \in H_i(\bar{s}_i), \forall s_i \in S_i(h_i), \quad U_i(\bar{s}_i, \mu^i(\cdot|h_i)) \geq U_i(s_i, \mu^i(\cdot|h_i)).$$

By Remark 8, weak sequential optimality is invariant under behavioral equivalence; therefore, it is a property of reduced strategies:

Remark 10. A strategy \bar{s}_i is weakly sequentially optimal under a belief system μ^i if and only if every $s_i \equiv_i \bar{s}_i$ is also weakly sequentially optimal under μ^i .

Like sequential optimality, also weak sequential optimality may be unsatisfiable if μ^i is not forward consistent.¹⁶ Focusing on forward consistent belief systems, weak sequential

¹⁶This is shown by Example 3, noticing that sequential and weakly sequential optimality coincide for strategies of Isa that choose c' at the root of Γ''' .

optimality can be given the following “*forward-planning*” interpretation. Suppose player i forms her beliefs according to $\mu^i \in \Delta_{\mathbb{F}}^{H_i}(S_{-i})$ as the play unfolds, although she may not know how she would revise her beliefs upon observing unexpected information, i.e., she may not be fully introspective. At the root, she plans to follow a (possibly partial) strategy that maximizes her expected payoff restricted to the collection of information sets $H_i(\mu^i|\emptyset)$ she deems possible under $\mu^i(\cdot|\emptyset)$. Let $\bar{s}_{i,\emptyset}$ be such a strategy. One way to compute $\bar{s}_{i,\emptyset}$ is to perform a folding-back planning algorithm on the restricted collection $H_i(\mu^i|\emptyset)$, that is, starting from information sets that are terminal within $H_i(\mu^i|\emptyset)$, and then focusing on behavior in the subcollection of information sets in $H_i(\mu^i|\emptyset)$ that are possible if such partial plan is implemented; indeed, by Remark 8, the specification of $\bar{s}_{i,\emptyset}$ at information sets in $H_i(\mu^i|\emptyset)$ that cannot occur under $\bar{s}_{i,\emptyset}$ is immaterial for the maximization problem. Since conditional beliefs within $H_i(\mu^i|\emptyset)$ are obtained by updating, player i would have no incentive to deviate from $\bar{s}_{i,\emptyset}$ at any $h_i \in H_i(\mu^i|\emptyset)$.

Example 6. Consider again game Γ''' and let μ^i be a forward consistent belief system of Isa such that $\mu^i(\ell|\emptyset) < \frac{1}{2}$. Then it is optimal to go down immediately, i.e., $\bar{s}_{a,\emptyset}(\emptyset) = d'$, (c') is inconsistent with $\bar{s}_{a,\emptyset}$, and the specification of $\bar{s}_{i,\emptyset}$ at (c') does not matter. \blacktriangle

If, implementing $\bar{s}_{i,\emptyset}$, player i unexpectedly obtains information h_i , where h_i is a first follower in $H_i \setminus H_i(\mu^i|\emptyset)$ of the information sets in $H_i(\mu^i|\emptyset)$, then i comes up with a new belief $\mu^i(\cdot|h_i) \in \Delta(S_{-i}(h_i))$ and plans to follow a (possibly partial) continuation strategy \bar{s}_{i,h_i} within the collection $H_i(\mu^i|h_i)$ of information sets weakly following h_i that she deems possible under $\mu^i(\cdot|h_i)$. If a hypothetical external observer knew μ^i (and how i breaks ties), she could determine \bar{s}_i for all information sets consistent with \bar{s}_i , that is, those in $H_i(\bar{s}_i)$, and \bar{s}_i would be weakly sequentially optimal under μ^i .¹⁷

The following result shows that the two notions of planning—weak sequential optimality and sequential optimality—are behaviorally equivalent. By Remark ??, behavioral and realization equivalence coincide. Hence, when a co-player j assesses the likely consequences of his choices by assigning subjective probabilities to the “ways of behaving” of a supposedly rational player i , the difference between the strong and weak version of optimality for i is immaterial. In other words, weak sequential optimality is sufficient to obtain the behavioral implications of rationality that are relevant for strategic reasoning (cf. Section 7).

Lemma 3. *A strategy is weakly sequentially optimal under a forward consistent belief system μ^i if and only if it is behaviorally equivalent to some strategy that is sequentially*

¹⁷The strategies obtained by “folding-back” on the collections $H_i(\mu^i|h_i)$ of information sets deemed possible by beliefs $\mu^i(\cdot|h_i)$ ($h_i \in H_i$) form a subset of $BR_i(\mu^i)$, those weakly sequentially optimal under μ^i (and a superset of $BR_i^*(\mu^i)$). But, for each one of them, there is a structurally equivalent strategy in $BR_i(\mu^i)$. We omit the details.

optimal under μ^i .

Proof. Suppose that \bar{s}_i is weakly sequentially optimal under $\mu^i \in \Delta_{\mathbb{F}^i}^{H_i}(S_{-i})$. By Corollary 2, there is some strategy \hat{s}_i that is sequentially optimal under μ^i . Let \bar{s}'_i denote the strategy that coincides with \bar{s}_i on $H_i(\bar{s}_i)$ and with \hat{s}_i on $H_i \setminus H_i(\bar{s}_i)$. By construction, \bar{s}'_i is behaviorally equivalent to \bar{s}_i . We must show that \bar{s}'_i is sequentially optimal under μ^i .

Since \bar{s}_i is weakly sequentially optimal under μ^i , Remark 8 implies that

$$\forall h_i \in H_i(\bar{s}_i), \forall s_i \in S_i(h_i), \quad U_i(\bar{s}'_i, \mu^i(\cdot|h_i)) = U_i(\bar{s}_i, \mu^i(\cdot|h_i)) \geq U_i(s_i, \mu^i(\cdot|h_i)).$$

Since \hat{s}_i is sequentially optimal under μ^i and \bar{s}'_i coincides with \hat{s}_i on $H_i \setminus H_i(\bar{s}_i)$,

$$\forall h_i \in H_i \setminus H_i(\bar{s}_i), \forall s_i \in S_i(h_i), \quad U_i(\rho_i[\bar{s}'_i/h_i], \mu^i(\cdot|h_i)) = U_i(\rho_i[\hat{s}_i/h_i], \mu^i(\cdot|h_i)) \geq U_i(s_i, \mu^i(\cdot|h_i)).$$

It follows that \bar{s}'_i is sequentially optimal under μ^i . \blacksquare

Clearly, the equivalence result also holds for standard and completely consistent belief systems. The following example illustrates Lemma 3.

Example 7. Strategy $R.a'.a''$ of Isa in game Γ' of Figure 1 is not sequentially optimal under any standard belief system, because there is no common value $\mu^i(C.\ell|(L, C)) = \mu^i(C.\ell|(R, C))$ that makes a' a best reply given (L, C) (which requires $\mu^i(C.\ell|(L, C)) \leq \frac{2}{5}$), and a'' a best reply given (R, C) (which requires $\mu^i(C.\ell|(R, C)) \geq \frac{1}{2}$). Yet, there is a continuum of standard belief systems that make the behaviorally equivalent strategy $R.d'.a''$ sequentially optimal, including all the $\mu^i \in \Delta^{H_i}(S_{-i})$ with $\mu^i(\{Q.\ell, Q.r\}|\emptyset) = 1$ and $\mu^i(C.\ell|(L, C)) = \mu^i(C.\ell|(R, C)) \geq \frac{1}{2}$ (there are more if $\varepsilon > 0$). More generally, it can be verified that every strategy s_i of Isa is weakly sequentially optimal under some standard belief system $\mu^i \in \Delta^{H_i}(S_{-i})$; hence, Lemma 3 implies that every s_i is behaviorally equivalent to some s'_i that is sequentially optimal under some $\mu^i \in \Delta^{H_i}(S_{-i})$. \blacktriangle

In the previous example, every strategy that can be justified as weakly sequentially optimal under some forward consistent belief system can also be justified as weakly sequentially optimal under some standard belief system, that is,

$$\cup_{\mu^i \in \Delta_{\mathbb{F}^i}^{H_i}(S_{-i})} BR_i(\mu^i) = \cup_{\mu^i \in \Delta^{H_i}(S_{-i})} BR_i(\mu^i).$$

Indeed, we prove in Section 7 that this is a general result. Yet, the following example shows

that the range of the weak sequential optimality map

$$\begin{aligned} BR_i : \Delta_{\mathbb{F}}^{H_i}(S_{-i}) &\rightarrow 2^{S_i} \\ \mu^i &\mapsto BR_i(\mu^i) \end{aligned}$$

may become strictly smaller when we restrict its domain by considering only standard belief systems. In other words, we may have some forward consistent belief system μ^i such that $BR_i(\mu^i) \neq BR_i(\bar{\mu}^i)$ for every standard belief system $\bar{\mu}^i$.

Example 8. Consider game Γ' of Figure 1 with $\varepsilon = 0$. Let

$$\Delta_{a'.a''}^o = \{\mu^i \in \Delta_{\mathbb{F}}^{H_i}(S_j) : BR_i(\mu^i) = L.a' \cup R.a''\}$$

denote the set of forward consistent belief systems of Isa justifying the reduced strategies $L.a'$ and $R.a''$ as weakly sequentially optimal. It can be verified that $\mu^i \in \Delta_{a'.a''}^o$ if and only if $\mu^i(\{Q.\ell, Q.r\}|\emptyset) = 1$, $\mu^i(C.\ell|(L, C)) < \frac{2}{5}$ and $\mu^i(C.\ell|(R, C)) > \frac{1}{2}$ (thus, $\Delta_{a'.a''}^o$ is the relative interior of subset $\Delta_{a'.a''}$ defined in Example 5). Indeed, every belief system that assigns probability 1 to Q is trivially forward consistent; if $\mu^i(\{Q.\ell, Q.r\}|\emptyset) < 1$ then $\mu^i(C.\ell|(L, C)) = \mu^i(C.\ell|(R, C))$, contradicting at least one of the aforementioned inequalities, which are those that make a' the unique best reply given (L, C) and a'' the unique best reply given (R, C) . Thus, none of the belief systems in $\Delta_{a'.a''}^o$ is standard: $\Delta_{a'.a''}^o \cap \Delta^{H_i}(S_j) = \emptyset$. Fix $\mu^i \in \Delta_{a'.a''}^o$ and $s_i \in BR_i(\mu^i)$. If $s_i \in L.a'$, then we can modify μ^i at $\{(R, C)\} \notin H_i(s_i)$ to obtain a standard belief system $\bar{\mu}^i$ with $\bar{\mu}^i(\{Q.\ell, Q.r\}|\emptyset) = \mu^i(\{Q.\ell, Q.r\}|\emptyset) = 1$ and

$$\bar{\mu}^i(C.\ell|(R, C)) = \bar{\mu}^i(C.\ell|(L, C)) = \mu^i(C.\ell|(L, C)) < \frac{2}{5}.$$

Then $s_i \in BR_i(\bar{\mu}^i) = L.a' \cup R.d''$. If instead $s_i \in R.a''$, then we can modify μ^i at $\{(L, C)\}$ to obtain a standard belief system $\bar{\mu}^i$ with $\bar{\mu}^i(\{Q.\ell, Q.r\}|\emptyset) = \mu^i(\{Q.\ell, Q.r\}|\emptyset) = 1$ and

$$\bar{\mu}^i(C.\ell|(L, C)) = \bar{\mu}^i(C.\ell|(R, C)) = \mu^i(C.\ell|(R, C)) > \frac{1}{2}.$$

Then $s_i \in BR_i(\bar{\mu}^i) = L.d' \cup R.a''$. Yet, for every $s_i \in BR_i(\mu^i) = L.a' \cup R.a''$ and every standard belief system $\bar{\mu}^i$ that coincides with μ^i on $H_i(s_i)$, $BR_i(\bar{\mu}^i) \neq L.a' \cup R.a''$. This implies that it is impossible to find a standard belief system $\bar{\mu}^i$ such that $BR_i(\bar{\mu}^i) = BR_i(\mu^i)$.

▲

To summarize, the analysis and discussion of rational planning in this section clarifies two points:

1. In order to rationally decide what to do at any information set h_i , player i only has to

know her current belief at h_i and what she would believe at followers she deems possible given h_i , that is, information sets $h'_i \in H_i(\mu^i|h_i)$. The resulting plan is dynamically consistent (optimal at each $h'_i \in H_i(\mu^i|h_i)$, or at least the elements of $H_i(\mu^i|h_i)$ possible under the plan itself) as long as the forward chain rule holds. Therefore, *partial introspection and the forward chain rule embedded in forward consistent belief systems are sufficient for rational planning* (even though the range of the weak sequential optimality map is affected by additional consistency requirements, as Example 8 shows).

2. *Forward planning* (iterated whenever a player is surprised, as explained above) *determines reduced rather than full strategies*, as anticipated in Section 2.

7 Behavioral implications of rationality and strategic reasoning

A **rational** player i endowed with a system of beliefs μ^i (which satisfies some posited consistency property) plays a strategy that is optimal given μ^i . In Section 6 we argued that the optimality criterion—sequential optimality obtained by folding back planning, or weak sequential optimality obtained by forward planning—does not matter if one is only interested in the observable features of the optimal strategy. Specifically, let $[s_i]_i \in S_i| \equiv_i$ denote the reduced strategy obtained from s_i .¹⁸ By Lemma 3, for any forward consistent belief system μ^i ,

$$BR_i(\mu^i) = \cup_{s_i^* \in BR_i^*(\mu^i)} [s_i^*]_i,$$

that is, the set of strategies that are weakly sequentially optimal under μ^i coincides with the set of strategies s_i that are behaviorally equivalent to some sequentially optimal strategy $s_i^* \in BR_i^*(\mu^i)$. Since behavioral and realization equivalence coincide (Remark 1), from the point of view of either an external observer who is interested in predicting the path of play, or a co-player who has to assess the likely consequences of her own behavior, the question is: What strategies of i (or, equivalently, what reduced strategies in $S_i| \equiv_i$) can be justified as weak sequential best replies to some system of beliefs that satisfies the posited consistency properties? Or, in other words, what are the behavioral implications of i 's rationality?

In this section, we show that such behavioral implications do not depend on whether the justifying belief system is just forward consistent or it satisfies the more demanding properties that hold when it is derived from a conditional probability system. An induction argument can then be used to show that the behavioral implications of rationality and “common belief in rationality” do not depend on the posited consistency properties. We focus

¹⁸Recall that $S_i| \equiv_i$ is the quotient set of S_i given the behavioral equivalence relation \equiv_i .

on strong rationalizability, a solution procedure that captures forward induction reasoning in sequential games. We show that the rationalizability procedures defined with completely consistent, standard, and forward consistent belief systems are equivalent, and we provide a characterization of such procedures that sheds light on why such equivalence holds.¹⁹ To make the paper more self-contained, the aforementioned equivalence result is preceded by a gentle introduction to some key ideas from the literature on strategic reasoning and solution concepts in sequential games.

7.1 Justifiability and strong rationalizability

The idea of “common belief in rationality” has been extended from simultaneous-move games to sequential games in different ways, depending on how players are assumed to revise their beliefs after unexpected moves of the co-players. In this section, we analyze strong rationalizability,²⁰ which is based on the idea that players “strongly believe” in the rationality and strategic sophistication of co-players. A player strongly believes an event if she would be always certain of that event unless she observed evidence in direct contradiction with it. Strong rationalizability characterizes the behavioral implications of *rationality and common strong belief in rationality*, which in turn is a formalization of the “best rationalization principle”: players always ascribe to their co-players the highest degree of strategic sophistication consistent with their past behavior, even when surprised by co-players’ behavior. This form of reasoning is also often referred to as “forward induction.” See Battigalli (1996) and Battigalli & Siniscalchi (2002).

Let it be transparent that each player i ’s belief system belongs to the set Δ_i (e.g., $\Delta_i = \Delta_{\mathbb{F}}^{H_i}(S_{-i})$), and let $\Delta = (\Delta_i)_{i \in I}$. As argued above, if player i is rational and has beliefs in Δ_i , then she may play any strategy s_i that can be justified as a weak sequential best reply to *some* belief system $\mu^i \in \Delta_i$. Thus, as a first step, the behavioral implications of i ’s rationality given Δ are characterized by the set

$$S_i^{\Delta,1} := \{s_i \in S_i : \exists \mu^i \in \Delta_i, s_i \in BR_i(\mu^i)\}.$$

We say that the strategies in $BR_i(\mu^i)$ are **justified by** μ^i , and that the strategies in $S_i^{\Delta,1}$ are **justifiable** for player i (relative to Δ).

To model the behavioral implications of rationality *and strategic reasoning*, as a second step we assume that each player “strongly believes” that the co-players play justifiable strate-

¹⁹Similar results hold for other versions of the rationalizability idea, see the discussion in Section 8.6.

²⁰Strong Rationalizability used to be called “extensive-form rationalizability” (Battigalli 1997). We avoid this terminology because there are different ways to formalize the rationalizability idea in the extensive-form analysis of sequential games. See Section 8.6.

gies, which follows from strong belief in the co-players' rationality (Battigalli & Siniscalchi 2002). Informally, an agent strongly believes an event E if she is initially certain of E and *continues to be certain of E* unless her information contradicts E . Since we only analyze the *behavioral implications* of rationality and strategic reasoning, here we only consider beliefs about co-players' behavior.²¹ Formally, a belief system μ^i **strongly believes** $E_{-i} \subseteq S_{-i}$, denoted $\mu^i \in \text{SB}_i(E_{-i})$, if $\mu^i(E_{-i}|h_i) = 1$ for all $h_i \in H_i$ such that $S_{-i}(h_i) \cap E_{-i} \neq \emptyset$. Strong belief in justifiability shapes how players revise their beliefs when they are surprised by their co-players' past behavior so as to capture the *rationalization principle*: a player should always try to interpret her information about co-players' past behavior by assuming they are playing strategies consistent with their subjective rationality. Thus, if i is rational and strongly believes in her opponents' rationality, she may play any strategy that can be justified by a belief system in Δ_i that strongly believes $S_{-i}^{\Delta,1}$, that is, any strategy in the set

$$S_i^{\Delta,2} := \{s_i \in S_i : \exists \mu^i \in \Delta_i \cap \text{SB}_i(S_{-i}^{\Delta,1}), s_i \in \text{BR}_i(\mu^i)\}.$$

Now the third step allows to point out an important feature of the notion of strategic reasoning captured by strong rationalizability. If i strongly believes $S_{-i}^{\Delta,2}$, she strongly believes that co-players' behavior is compatible with one degree of strategic sophistication, rationality and strong belief in others' rationality; yet it may be that an information set h_i contradicts such degree of sophistication, that is, $S_{-i}(h_i) \cap S_{-i}^{\Delta,2} = \emptyset$. In this case strong belief in $S_{-i}^{\Delta,2}$ leaves i 's conditional beliefs at h_i unrestricted, so i may well assign positive probability at h_i to some *unjustifiable* strategies of the co-players, *even if h_i is compatible with justifiability* (hence, with $-i$ ' rationality), thus violating the rationalization principle. In particular, this can happen if $S_{-i}^{\Delta,2} \subset S_{-i}^{\Delta,1}$ and $S_{-i}(h_i) \cap S_{-i}^{\Delta,1} \neq \emptyset$. Therefore, we shall require strong belief in *both* $S_{-i}^{\Delta,2}$ and $S_{-i}^{\Delta,1}$, i.e., $\mu^i \in \text{SB}_i(S_{-i}^{\Delta,1}) \cap \text{SB}_i(S_{-i}^{\Delta,2})$ (besides $\mu^i \in \Delta_i$). This captures the *best rationalization principle*. Continuing further this reasoning process, we capture higher degrees of strategic sophistication and, according to the posited notion of belief system, we obtain the definition of strong rationalizability. To express this formally, given a set of belief systems Δ_i and events $E_{-i}^1, \dots, E_{-i}^n \subseteq S_{-i}$, let

$$J_i(\Delta_i; E_{-i}^1, \dots, E_{-i}^n) := \{s_i \in S_i : \exists \mu^i \in \Delta_i \cap (\bigcap_{k=1}^n \text{SB}_i(E_{-i}^k)), s_i \in \text{BR}_i(\mu^i)\}$$

denote the set of **strategies justifiable by belief systems in Δ_i that strongly believe $E_{-i}^1, \dots, E_{-i}^n$** . Thus, for example, $J_i(\Delta_i; S_{-i}) = S_i^{\Delta,1}$, because S_{-i} is trivially strongly believed by every μ^i , and $J_i(\Delta_i; S_{-i}^{\Delta,1}) = S_i^{\Delta,2}$.

²¹That is, systems of **first-order beliefs**: see, e.g., Battigalli & Siniscalchi (2002) for the connection to higher-order beliefs and to strong belief in the co-players' rationality and strategic sophistication.

Definition 12. Fix $\Delta = (\Delta_i)_{i \in I} \in \left\{ (\Delta_F^{H_i}(S_{-i}))_{i \in I}, (\Delta^{H_i}(S_{-i}))_{i \in I}, (\Delta_C^{H_i}(S_{-i}))_{i \in I} \right\}$.

Let $(S_i^{\Delta,1})_{i \in I} := (J_i(\Delta_i; S_{-i}))_{i \in I}$ and define recursively, for each $m \in \mathbb{N}$ and $i \in I$, $S_i^{\Delta,m+1} = J_i(\Delta_i; S_{-i}^{\Delta,1}, \dots, S_{-i}^{\Delta,m})$. The set of **strongly rationalizable** strategies for i **with forward consistent** (respectively, **standard**, **completely consistent**) **belief systems** is $S_i^{\Delta,\infty} := \bigcap_{k=1}^{\infty} S_i^{\Delta,k}$ where $\Delta = (\Delta_F^{H_i}(S_{-i}))_{i \in I}$ (respectively, $\Delta = (\Delta^{H_i}(S_{-i}))_{i \in I}$, $\Delta = (\Delta_C^{H_i}(S_{-i}))_{i \in I}$).

The best rationalization principle is captured by strong belief in each event of the sequence $(S_{-i}^{\Delta,m})_{m=0}^{\infty}$. To see this, notice first that $(S_{-i}^{\Delta,k})_{k=0}^{\infty}$ is a (weakly) decreasing sequence of subsets, since each step adds additional (strong-belief) restrictions on the set of beliefs systems allowed to justify a strategy. Since S_{-i} is finite, the sequence must become constant after some finite number of steps K . Then, letting $\kappa(h_i)$ denote the highest degree of strategic sophistication compatible with h_i ,²² $\mu^i \in \bigcap_{k=0}^{\infty} \text{SB}_i(S_{-i}^{\Delta,k})$ implies $\mu^i(S_{-i}^{\Delta,\kappa(h_i)} | h_i) = 1$ for each $h_i \in H_i$. Note also that, by Remark 10, strong rationalizability may be interpreted as an iterated elimination of reduced strategies.

7.2 Sufficiency of forward consistency

We are now ready for the main result of the paper: strong rationalizability is invariant to the use of forward consistent, standard, or completely consistent belief systems. We start with an example of this fact.

Example 9. Consider game Γ' with $\varepsilon > 0$. Every version of strong rationalizability solves the game in three steps. Indeed, for every notion of belief system, and focusing for simplicity on reduced strategies, we get

step	Isa	Joe
1	$\{L.a', L.d', R.a'', R.d''\}$	$\{Q, C.r\}$
2	$\{R.d''\}$	$\{Q, C.r\}$
3	$\{R.d''\}$	$\{C.r\}$

Step 1: If Joe Quits at the outset he gets at least 6 utils. If he Continues, he can get a payoff larger than 6 only if Isa plays $R.d''$, allowing him to get 8 utils if he then goes left at information set $h_j = \{((L, C), d'), ((R, C), d'')\}$. Hence, the optimal choice of Joe at the initial history depends on the probability he assigns to $R.d''$. If he assigns sufficiently high probability to $R.d''$, and then updates accordingly at h_j , it is optimal for him to play $C.r$, otherwise, it is optimal for him to play Q . Note that the sets of forward consistent, standard,

²²That is, $\kappa(h_i) = \max \left\{ k \in \{0, \dots, K\} : S_{-i}(h_i) \cap S_{-i}^{\Delta,k} \neq \emptyset \right\}$

and completely consistent belief systems of Joe all coincide, because his two information sets are ordered by precedence, therefore his beliefs are fully disciplined by the forward chain rule.

Step 2: If Isa is initially certain that Joe is rational, she must initially assign probability 1 to $\{Q, C.r\}$. If she assigns positive probability to $C.r$, her best reply is $R.d''$, and if she does play R and observes that Joe played C (information set $h_i'' = \{(R, C)\}$), her updated belief will assign probability 1 to $C.r$. Then, she will optimally go down, thus implementing $R.d''$. If instead she initially assigns probability 1 to Q , it is still optimal for her to play R at the root, but she would be surprised to observe that Joe continued. Nonetheless, since C is compatible with Joe's rationality, if she *strongly* believes that Joe is rational, also in this case at h_i'' she must assign probability 1 to $C.r$ and optimally go down, thus implementing $R.d''$. Therefore, under strong belief in Joe's rationality, only $R.d''$ is justifiable for Isa. Note that this behavior of Isa only depends on two beliefs: the initial belief, and the belief she would have after choosing her justifiable action R , and upon observing that Joe continued. The relationship between these two beliefs is disciplined by the forward chain rule under all three notions of belief system. If Isa has a standard belief system, the belief she would have after choosing L and observing that Joe continued, i.e., at information set $h_i' = \{(L, C)\}$, must be identical to the one at information h_i'' , because the two information sets convey the same information about Joe's behavior. However, Isa's belief at h_i' plays no role in determining her justifiable strategies, because she does not plan to choose L . For this reason, all notions of belief system induce the same behavior of Isa.

Step 3: If Joe is initially certain that Isa is rational and that she strongly believes in his rationality, he assigns probability 1 to $R.d''$ and plays $C.r$, for the reason illustrated at step 1. Thus, strong rationalizability pins down one (reduced) strategy for each player: $R.d''$ for Isa and $C.r$ for Joe. ▲

To determine Isa's justifiable behavior at step 2, only her initial belief and her belief at the information set that is consistent with her own initial move matter. What she would believe after an initial move that she had not planned to make is irrelevant for *weak* sequential optimality, hence is irrelevant to determine her justifiable behavior. Indeed, whether a strategy s_i is justified by belief system μ^i (that is, $s_i \in BR_i(\mu^i)$) only depends on the *partial* belief system $(\mu^i(\cdot|h_i))_{h_i \in H_i(s_i)}$. With this, for any fixed strategy s_i , how do the different notions of consistency affect the partial belief systems with domain $H_i(s_i)$? The key take-away from the example is that there is no difference between forward consistent and standard belief systems defined on this restricted domain. To see why, note first that for every strategy s_i , the collection of s_i -possible information sets $H_i(s_i)$ is such that *precedence between these information sets mirrors the accumulation of information about co-players' behavior*: given two information sets $h_i, h_i' \in H_i(s_i)$, either they are ordered by precedence, and then

the beliefs $\mu^i(\cdot|h_i)$ and $\mu^i(\cdot|h'_i)$ are disciplined by the forward chain rule, or they represent mutually exclusive observations of the co-players' past moves (i.e., $S_{-i}(h_i) \cap S_{-i}(h'_i) = \emptyset$), making $\mu^i(\cdot|h_i)$ and $\mu^i(\cdot|h'_i)$ unrelated. This is because, for each $s_{-i} \in S_{-i}$, the information sets of player i along the induced path $\zeta(s_i, s_{-i})$ are ordered by perfect recall. Intuitively, if i plays some given strategy s_i , it is impossible that the same moves of the co-players lead to two mutually exclusive information sets of i .

Actually, the fact that the forward chain rule disciplines the beliefs conditional on the information sets in $H_i(s_i)$ as much as possible allows to extend a partial belief system on $H_i(s_i)$ that satisfies the forward chain rule to a full-blown belief system on H_i that satisfies the strongest notion of consistency, i.e., complete consistency. Furthermore, such extension is possible even if we require strong belief in an event, or a chain of events, as we have to do when we consider steps 2, 3, ... of the strong rationalizability solution procedure. Formally, for any given strategy s_i , *partial* belief system $\mu^i \in \times_{h_i \in H_i(s_i)} \Delta(S_{-i}(h_i))$ is **forward consistent**—written $\mu^i \in \Delta_{\mathbf{F}}^{H_i(s_i)}(S_{-i})$ —if

$$\mu^i(S_{-i}(h'_i)|h_i) > 0 \Rightarrow \mu^i(E_{-i}|h'_i) = \frac{\mu^i(E_{-i}|h_i)}{\mu^i(S_{-i}(h'_i)|h_i)}$$

for all $h_i, h'_i \in H_i(s_i)$ with $h_i \prec_i h'_i$, and all $E_{-i} \subseteq S_{-i}(h'_i)$, and μ^i **strongly believes** an event $E_{-i}^* \subseteq S_{-i}$ —written $\mu^i \in \text{SB}_i(\mu^i)$ —if $\mu^i(E_{-i}^*|h_i) = 1$ for all $h_i \in H_i(s_i)$ such that $E_{-i}^* \cap S_{-i}(h_i) \neq \emptyset$.

Given our interest in the best rationalization principle, we consider belief systems that strongly believe each element of a **decreasing chain** of events $\mathcal{E}_{-i} = (E_{-i}^1, \dots, E_{-i}^n)$, where $\emptyset \neq E_{-i}^n \subseteq \dots \subseteq E_{-i}^1 \subseteq S_{-i}$.

Lemma 4. *Fix a strategy s_i and a forward consistent belief system μ^i on $H_i(s_i)$ that strongly believes each element of a decreasing chain $\mathcal{E}_{-i} = (E_{-i}^1, \dots, E_{-i}^n)$; then there is a completely consistent belief system $\bar{\mu}^i$ that strongly believes each element of \mathcal{E}_{-i} such that $\bar{\mu}^i(\cdot|h_i) = \mu^i(\cdot|h_i)$ for all $h_i \in H_i(s_i)$.*

Proof: CLAIM 1: Collection $\mathcal{H}_i(s_i) := \{S_{-i}(h_i) : h_i \in H_i(s_i)\}$ is an event tree. To see this, first note that $S_{-i} = S_{-i}(\{\emptyset\}) \in \mathcal{H}_i(s_i)$. Next fix $h_i, \bar{h}_i \in H_i(s_i)$ arbitrarily and suppose there exists $s_{-i}^* \in S_{-i}(h_i) \cap S_{-i}(\bar{h}_i)$. We must show that either $S_{-i}(\bar{h}_i) \subseteq S_{-i}(h_i)$, or $S_{-i}(h_i) \subseteq S_{-i}(\bar{h}_i)$. By assumption, $s_i \in S_i(h_i)$ and $s_i \in S_i(\bar{h}_i)$. Then, $(s_i, s_{-i}^*) \in S_i(h_i) \times S_{-i}(h_i)$ and $(s_i, s_{-i}^*) \in S_i(\bar{h}_i) \times S_{-i}(\bar{h}_i)$. By perfect recall, $S(h) = S_i(h) \times S_{-i}(h)$ for each $h \in H_i$. Therefore, $(s_i, s_{-i}^*) \in S(h_i) \cap S(\bar{h}_i)$. Since $S(h_i) \cap S(\bar{h}_i) \neq \emptyset$, by perfect recall, either $S(\bar{h}_i) \subseteq S(h_i)$ and $S_{-i}(\bar{h}_i) \subseteq S_{-i}(h_i)$, or $S(h_i) \subseteq S(\bar{h}_i)$ and $S_{-i}(h_i) \subseteq S_{-i}(\bar{h}_i)$.

CLAIM 2: Let μ^i be as in the statement. Then, there is a CPS $\hat{\mu}^i \in \Delta^{\mathcal{H}_i(s_i)}(S_{-i})$ that strongly believes each element of \mathcal{E}_{-i} and is such that $\mu^i(\cdot|h_i) = \hat{\mu}^i(\cdot|S_{-i}(h_i))$ for all

$h_i \in H_i(s_i)$. Intuitively, this follows from the fact that the information sets in $H_i(s_i)$ mirror the accumulation of information about co-players. For a formal proof see the Appendix.

Let $\hat{\mu}^i \in \Delta^{H_i(s_i)}(S_{-i})$ be as per CLAIM 2. Lemma 1 implies that there is a complete CPS $\widehat{\mu}^i \in \Delta^*(S_{-i})$ that extends $\hat{\mu}^i$ and strongly believes each element of \mathcal{E}_{-i} . Let $\bar{\mu}^i(\cdot|h_i) = \widehat{\mu}^i(\cdot|S_{-i}(h_i))$ for all $h_i \in H_i$. By construction, $\bar{\mu}^i$ is a completely consistent belief system that strongly believes each element of \mathcal{E}_{-i} and coincides with partial belief system μ^i on $H_i(s_i)$. ■

By using partial belief systems on $H_i(s_i)$, Lemma 4 offers an operationally useful characterization of justifiability and yields our main results. First note that, for any strategy s_i and *partial* belief system $\mu^i \in \times_{h_i \in H_i(s_i)} \Delta(S_{-i}(h_i))$, it makes sense to say that s_i is justified by μ^i , written $s_i \in BR_i(\mu^i)$, if

$$\forall h_i \in H_i(s_i), \forall s'_i \in S_i(h_i), \quad U_i(s_i, \mu^i(\cdot|h_i)) \geq U_i(s'_i, \mu^i(\cdot|h_i)).$$

With this, we can define the set of strategies of player i **justifiable by partial**, forward consistent belief systems that strongly believe events $E_{-i}^1, \dots, E_{-i}^n$:

$$J_i^p(E_{-i}^1, \dots, E_{-i}^n) := \left\{ s_i \in S_i : \exists \mu^i \in \Delta_{\mathbb{F}}^{H_i(s_i)}(S_{-i}) \cap \left(\bigcap_{m=1}^n \text{SB}_i(E_{-i}^m) \right), s_i \in BR_i(\mu^i) \right\}.$$

Lemma 4 implies the following key result.

Theorem 4. *Fix a decreasing chain $\mathcal{E}_{-i} = (E_{-i}^1, \dots, E_{-i}^n)$. A strategy $s_i \in S_i$ is justifiable by a completely consistent, or standard, or forward consistent belief system that strongly believes each element of \mathcal{E}_{-i} if and only if it is justifiable by some partial, forward consistent belief system on $H_i(s_i)$ that strongly believes each element of \mathcal{E}_{-i} ; that is,*

$$\begin{aligned} J_i(\Delta_C^{H_i}(S_{-i}); E_{-i}^1, \dots, E_{-i}^n) &= J_i(\Delta^{H_i}(S_{-i}); E_{-i}^1, \dots, E_{-i}^n) = J_i(\Delta_{\mathbb{F}}^{H_i}(S_{-i}); E_{-i}^1, \dots, E_{-i}^n) \\ &= J_i^p(E_{-i}^1, \dots, E_{-i}^n). \end{aligned}$$

Theorem 4 suggests a simplified rationalizability algorithm whereby strategies are justified by partial belief systems. For each $i \in I$, let

$$S_i^{p,1} := J_i^p(S_{-i}) = \left\{ s_i \in S_i : \exists \mu^i \in \Delta_{\mathbb{F}}^{H_i(s_i)}(S_{-i}), s_i \in BR_i(\mu^i) \right\}$$

denote the strategies that are justified by *partial*, forward consistent belief systems. By Theorem 4, $S_i^{p,1}$ is also the set of justifiable strategies of i for each one of the three notions

of “total” consistent belief system. For each $m \in \mathbb{N}$, recursively define

$$S_i^{p,m+1} := J_i^p (S_{-i}^1, \dots, S_{-i}^m).$$

Given Theorem 4, a straightforward induction argument yields the main result:

Theorem 5. *The sets of strongly rationalizable strategies with forward consistent, standard, and completely consistent belief systems coincide; in particular,*

$$(S^{\Delta,m})_{m=1}^\infty = (S^{p,m})_{m=1}^\infty$$

for each $\Delta \in \left\{ (\Delta_{\text{F}}^{H_i}(S_{-i}))_{i \in I}, (\Delta^{H_i}(S_{-i}))_{i \in I}, (\Delta_{\text{C}}^{H_i}(S_{-i}))_{i \in I} \right\}$.

The aforementioned definitions of strong rationalizability require that, at step $m + 1$, the justifying belief system strongly believes the behavioral events corresponding to each previous step $1, \dots, m$. While this transparently captures the best rationalization principle, it adds a layer of complexity to the solution procedure. Yet, building on previous work on strong rationalizability (cf. Battigalli 1997), one can characterize it with a simpler reduction algorithm that only requires, at step $m + 1$, strong belief in the set of co-players’ strategies that survived up to step m . We omit the details.

8 Discussion

In this section we discuss some features and extensions of our analysis, and how it relates to the extant literature.

8.1 A comment on partial belief systems

In Section 7, we derived our main equivalence results by looking at partial belief systems defined on the subcollection of information sets of a player that are possible when she plays a given strategy. We emphasize that partial belief systems are just an analytical tool. The conceptually primary belief systems are the “total” ones considered in the analysis of rational planning of Section 6. Indeed, when one reasons strategically about player i , the justifiability of a strategy s_i by a partial belief system on $H_i(s_i)$ does not represent how the analyst or a co-player thinks about i ’s planning, it is only a step in verifying whether the behavior described by s_i (or any equivalent strategy) is consistent with rationality. Since the precedence relation in $H_i(s_i)$ mirrors the accumulation of information about the co-players’ behavior, for such partial belief systems forward consistency is equivalent to stronger versions of the chain rule. Thus, the fact that only such partial belief systems matter gives insights for why the

behavioral implications of rationality and strategic reasoning do not depend on consistency properties beyond forward consistency.

8.2 Beliefs at information sets where players are active

In the seminal works of von Neumann & Morgenstern (1944) and Kuhn (1953), and in most of the following game theoretic work, the information of players is specified only at information sets where they are active. Indeed, it was either argued or taken as self-evident that the information of inactive players is irrelevant for the strategic analysis of games. We can confirm, to some extent, this intuition.²³

Let $\hat{H}_i := \{h_i \in H_i : |\mathcal{A}_i(h_i)| > 1\} \cup \{\emptyset\}$ denote the collection of information sets where a player is active plus the root.²⁴ Considering this smaller collection does not change the characterization of weak sequential optimality: a strategy s_i is weakly sequentially optimal under a belief system μ^i if and only if it maximizes expected payoff conditional on each $h_i \in \hat{H}_i(s_i) := H_i(s_i) \cap \hat{H}_i$. Yet, by restricting the domain of belief systems, the chain rule loses bite. Does this affect solution concepts? We show that, for strong rationalizability, the answer is negative.

As a preliminary observation, note that $\hat{H}_i(s_i)$ inherits from $H_i(s_i)$ the property that precedence in $\hat{H}_i(s_i)$ mirrors the accumulation of information about the co-players. With this, we can always extend a forward consistent belief system on $\hat{H}_i(s_i)$ to a completely consistent “total” one in a way that preserves strong belief in a chain of events. In fact, our proof for the extension of forward consistent belief systems on $H_i(s_i)$ works *without modification* when replacing $H_i(s_i)$ with $\hat{H}_i(s_i)$. Thus:

Lemma 5. *Fix a strategy s_i and a forward consistent belief system μ^i on $\hat{H}_i(s_i)$ that strongly believes each element of the decreasing chain $\mathcal{E}_{-i} = (E_{-i}^1, \dots, E_{-i}^n)$; then there is a completely consistent belief system $\bar{\mu}^i$ on H_i that strongly believes each element of \mathcal{E}_{-i} such that $\bar{\mu}^i(\cdot|h_i) = \mu^i(\cdot|h_i)$ for all $h_i \in \hat{H}_i(s_i)$.*

Thus, given the obvious definitions of justifiability by a (completely consistent, standard, forward consistent) belief system on \hat{H}_i and $\hat{H}_i(s_i)$, we obtain the following:

Theorem 6. *Fix a strategy s_i and a decreasing chain $\mathcal{E}_{-i} = (E_{-i}^1, \dots, E_{-i}^n)$. The following statements are equivalent:*

²³Such confirmation can only be partial. When some forms of psychological motivations are incorporated in the game theoretic framework, the information of inactive players may be very relevant. See, e.g., Battigalli & Dufwenberg (2022) and the relevant references therein.

²⁴Adding the root is just a technical convenience and it makes conceptual sense: after her strategic analysis of the game, player i starts with some belief $\mu^i(\cdot|\emptyset)$.

1. s_i is justifiable by a (completely consistent, standard, forward consistent) belief system on \hat{H}_i that strongly believes each element of \mathcal{E}_{-i} ;
2. s_i is justifiable by a (completely consistent, standard, forward consistent) belief system on H_i that strongly believes each element of \mathcal{E}_{-i} ;
3. s_i is weakly sequentially optimal under some partial, forward consistent belief system on $H_i(s_i)$ that strongly believes each element of \mathcal{E}_{-i} ;
4. s_i is weakly sequentially optimal under some partial, forward consistent belief system on $\hat{H}_i(s_i)$ that strongly believes each element of \mathcal{E}_{-i} .

The belief systems on $\hat{H}_i(s_i)$ are even simpler objects to work with for the justifiability analysis than those on $H_i(s_i)$. For example, for a player that is active at most once along any path of play, for every s_i we have $H_i(s_i) = H_i$ and $\hat{H}_i(s_i) = \hat{H}_i$. But in \hat{H}_i there cannot be information sets related by precedence, and so forward consistency has no bite for belief systems on \hat{H}_i , and one can work with belief systems without worrying about any form of coherence of beliefs at different information sets. Given the obvious definitions of strong rationalizability with belief systems on \hat{H}_i for each i , Theorem 6 and a straightforward induction argument yield:

Theorem 7. *The sets of strongly rationalizable strategies with forward consistent, standard, and completely consistent belief systems on information sets where players are active coincide, and they are the same as those obtained with belief systems on all information sets.*

8.3 Games with incomplete information

Our analysis extends seamlessly to finite games with incomplete information. Without essential loss of generality, incomplete information can be represented by assuming that the payoffs of some (or all) terminal nodes depend on a profile of parameters $\theta = (\theta_i)_{i \in I}$ in a finite set $\Theta = \times_{i \in I} \Theta_i$, so that each player i only knows his type θ_i (see, e.g., Battigalli & Siniscalchi 1999, 2002). With this, the previous analysis applies to an auxiliary game obtained by letting a fictitious and indifferent player choose θ and redefining information sets to take into account private information about θ : for each $h_i \in H_i$ in the original game and every type θ_i consider the information set

$$[\theta_i, h_i] = \{(\theta', x) \in \Theta \times X : \theta'_i = \theta_i, x \in h_i\}.$$

Justifiability and rationalizability in this auxiliary game correspond to justifiability and rationalizability in the game with incomplete information.

8.4 Directed rationalizability

Strong directed rationalizability posits, for each player i , a restricted subset of belief systems $\bar{\Delta}_i$. Each profile $\bar{\Delta} = (\bar{\Delta}_i)_{i \in I}$ yields a corresponding definition of strong $\bar{\Delta}$ -rationalizability as in Section 7, with $\Delta_{\mathbb{F}}^{H_i}(S_{-i})$, $\Delta^{H_i}(S_{-i})$, or $\Delta_{\mathbb{C}}^{H_i}(S_{-i})$ replaced by $\bar{\Delta}_i \subseteq \Delta_{\mathbb{F}}^{H_i}(S_{-i})$, $\bar{\Delta}_i \subseteq \Delta^{H_i}(S_{-i})$, or $\bar{\Delta}_i \subseteq \Delta_{\mathbb{C}}^{H_i}(S_{-i})$ (cf. Battigalli & Siniscalchi 2003).²⁵ The construction in the proof of our extension Lemma 4 preserves membership to restricted subsets of belief systems whenever these are obtained by imposing that *initial beliefs* satisfy some properties, that is, $\mu^i(\cdot|\emptyset) \in \bar{\Delta}_{\emptyset,i} \subseteq \Delta(S_{-i})$. Therefore, *our analysis extends to strong directed rationalizability as long as one considers only restrictions on initial beliefs*: in this case, the assumed consistency properties do not affect the solution. On one hand, restrictions on initial beliefs play an important role for several reasons. (i) In an incomplete-information scenario, the analyst may wish to just posit restrictions on players' **exogenous** beliefs, i.e., their *initial* beliefs about the asymmetrically known features θ of the game (cf. Section 8.3). In this case, strong $\bar{\Delta}$ -rationalizability is non-empty (Battigalli & Siniscalchi 2003) and the set of strongly rationalizable paths is monotone with respect to the posited restrictions,²⁶ whereas neither result holds for general restrictions involving endogenous beliefs.²⁷ (ii) The Iterated Intuitive Criterion is characterized by (the non-emptiness of) a version of strong directed rationalizability for signaling games that only posits restrictions on initial beliefs (Battigalli & Siniscalchi 2002, 2003). (iii) Similar considerations apply to the characterization of path-agreements that are self-enforcing under forward-induction reasoning (Catonini 2021). (iv) Finally, one can show that when $\bar{\Delta}$ only restricts initial beliefs, strong $\bar{\Delta}$ -rationalizability can be computed with a simplified “one-step-memory” algorithm (see our comment at the end of Section 7), while this is not true in general. On the other hand, when $\bar{\Delta}$ also features restrictions on *conditional* beliefs, the assumed notion of consistency may affect the result.

Example 10. Consider game Γ' with $\varepsilon = 0$ and the following restricted sets of (forward consistent, or standard) belief systems for Isa and Joe:

$$\begin{aligned}\bar{\Delta}_i &= \{ \mu^i : \mu^i(C.r|(L,C)) = 1 \}, \\ \bar{\Delta}_j &= \{ \mu^j : \mu^j(L.a'|\emptyset) = 1 \}.\end{aligned}$$

Allowing for all forward consistent belief systems that satisfy these restrictions, the set of strongly $\bar{\Delta}$ -rationalizable pairs of (reduced) strategies is $\{L.a', R.d'', R.a''\} \times \{Q\}$. In partic-

²⁵The solution set characterizes the behavioral implications of rationality and common strong belief in rationality under the assumption that the posited restrictions are transparent to the players (see, e.g., Battigalli & Friedenberg 2012).

²⁶Since strong belief is not monotone, this path-monotonicity result is not obvious. Its proof is available upon request.

²⁷That is, beliefs about behavior, or beliefs about exogenous features conditional on observed behavior.

ular, $R.a''$ is included because forward consistency allows Isa to assign a low probability to $C.r$ conditional on (R, C) if she is initially certain of Q , even though $\mu^i(C.r|(L, C)) = 1$. Standard belief systems, instead, satisfy $\mu^i(C.r|(L, C)) = \mu^i(C.r|(R, C))$, forcing Isa to assign probability 1 to $C.r$ conditional on (R, C) . Thus, the set of strongly $\bar{\Delta}$ -rationalizable pairs when one considers only standard belief systems is $\{L.a', R.d''\} \times \{Q\}$. \blacktriangle

8.5 Strong best-response sets

As in Section 7, let $\Delta = (\Delta_i)_{i \in I}$ denote the profile of spaces of forward consistent, or standard, or completely consistent belief systems. A Cartesian product of strategy subsets $K = \times_{i \in I} K_i$ is a **strong Δ -best response set (SBRs)** if, for every $i \in I$ and $s_i \in K_i$, there is a belief system $\mu^i \in \Delta_i \cap \text{SB}_i(K_{-i})$ such that $s_i \in \text{BR}_i(\mu^i) \subseteq K_i$. In words, each strategy in subset K_i must be justified as weakly sequentially optimal under some belief system in Δ_i that strongly believes K_{-i} and does *not* justify strategies outside K_i .²⁸ Battigalli & Friedenberg (2012) provide an epistemic justification of this concept,²⁹ showing that K is an SBRs if and only if it is the set of strategy profiles that are possible under rationality and common strong belief in rationality given some transparent contextual restrictions on belief hierarchies, that is, for some epistemic type structure. The set of strongly rationalizable strategy profiles (by Theorem 5, under any notion of belief system) is an SBRs, but it need not be the largest one, i.e., it may not contain some other SBRs, due to the non-monotonicity of strong belief. Thus, from the perspective of an analyst who does not know what is transparent to the players, the robust behavioral implications of rationality and common strong belief in rationality are characterized by the union of all the SBRs.

The following example shows that the union of all SBRs of a game is *not* invariant to the assumed notion of belief system. (Of course, this also shows that the *collection* of all SBRs is not invariant either.)

Example 11. Consider the variation of game Γ' (with $\varepsilon = 0$) depicted in Figure 11.

Compared to Γ' , Isa has an additional action (M) at the root, and Joe has an additional action (m) at h_j . Moreover, this is not a common interest game any longer, nor a multistage game, because at h_j Joe cannot distinguish between $L.d'$, $R.d''$, and M , so he does not know if Isa has already moved once or twice. In Figure 11, the payoffs of z are denoted by $(u_i(z), u_j(z))$ when they differ across players.

We are going to show that $C.m$ is part of a SBRs with forward consistency but not with standard consistency.

²⁸Battigalli & Friedenberg (2012) say “extensive-form best response set” and define them using standard belief systems.

²⁹By means of standard belief systems; see Section 8.7.




Figure 6: Modified running example.

First, let Δ denote the profile of spaces of forward consistent belief systems. We show that

$$K_i \times K_j = \{L.a', R.a''\} \times \{Q, C.l, C.m, C.r\}$$

is a strong Δ -best response set. Consider the forward consistent belief system μ^i of Isa such that

$$\mu^i(Q|\emptyset) = \mu^i(C.r|(L,C)) = \mu^i(C.l|(R,C)) = 1.$$

Note that μ^i strongly believes K_j . Moreover, since $BR_i(\mu^i) = K_i$, $K_i \times K_j$ satisfies the requirements of a SBRS for Isa. For Joe, fix first a belief system μ^j that strongly believes K_i such that $\mu^j(R.a''|\emptyset) = 1$ —since $\mu^j(S_i(h_j)|\emptyset) = 0$, μ^j is forward consistent. Since $BR_j(\mu^j) = \{Q\}$, Q satisfies the required property. Consider now a belief system μ^j that strongly believes K_i such that $\mu^j(L.a'|\emptyset) = 1$ —since $\mu^j(S_i(h_j)|\emptyset) = 0$, μ^j is forward consistent. At h_j , such a belief system can assign probability 1 to $L.d'$, and then $BR_j(\mu^j) = \{C.l\}$, or to $R.d''$, and then $BR_j(\mu^j) = \{C.r\}$, or to M , and then $BR_j(\mu^j) = C$. We conclude that $K_i \times K_j$ satisfies the requirements of a SBRS also for Joe.

Now let Δ denote the profile of spaces of standard belief systems. Consider a strong Δ -best response set $K_i \times K_j$. For Isa, strategy M is strictly dominated, therefore $M \notin K_i$. Moreover, every $\mu^i \in \Delta_i$ that justifies $L.a'$ also justifies $R.d''$: to justify $L.a'$, we need

$\mu^i(\{C.\ell, C.m\} | (L, C)) < 1/2$, hence $\mu^i(C.r | (L, C)) > 1/2$, and standard consistency imposes $\mu^i(\cdot | (R, C)) = \mu^i(\cdot | (L, C))$. Therefore, if $L.a' \in K_i$, then $R.d'' \in K_i$. For Joe, choosing C is justified only by belief systems that initially assign positive probability to $L.a'$, therefore either $K_j = \{Q\}$, or $L.a' \in K_i$. In the second case, since $M \notin K_i$ but $R.d'' \in K_i$, every belief system μ^j that strongly believes K_i cannot assign positive probability to M at h_j , which implies $C.m \notin BR_j(\mu^j)$. Therefore, $C.m \notin K_j$. \blacktriangle

8.6 Other versions of rationalizability

Strong rationalizability characterizes the behavioral implications of rationality and common strong belief in rationality, where the latter relies on the best rationalization principle. Theorem 5 shows that strong rationalizability is invariant to the adopted notion of consistency of belief systems. Other assumptions on strategic reasoning are worth considering and their behavioral implications are characterized by different versions of the rationalizability idea. *Initial rationalizability* obtains if it is only assumed that players have common belief in rationality at the beginning of the game, with no restrictions on how they revise their beliefs when surprised (Ben-Porath 1997, Battigalli & Siniscalchi 1999). Lemma 4 and Theorem 4 imply that the invariance result holds for this solution concept in a simplified way, because there are no strong belief conditions to preserve when one extends a partial forward consistent belief system to a total completely consistent belief system. *Backwards rationalizability* is more easily understandable for games with observable actions: it relies on the assumption that, for each non-terminal history, there is common belief in players' rationality in the continuation game even if someone previously made a "mistake" (Perea 2014, Penta 2015, Battigalli & De Vito 2021, Catonini & Penta 2022). In Appendix B we prove the invariance result for backwards rationalizability in games with observable actions.

8.7 Relation to previous work

On one hand, in most of the literature on rationalizability in sequential games, solution procedures are defined by considering forward consistent belief systems; see, e.g., Battigalli & Siniscalchi (2003) on strong Δ -rationalizability, or Perea (2014) on backwards rationalizability. A large body of the literature on strong rationalizability, such as Pearce (1984), Battigalli (1997), Shimoji & Watson (1998), and Shimoji (2004) use a notion of "consistent systems of conjectures" that is equivalent to forward consistent belief systems in terms of (weak) sequential optimality.³⁰ Complete CPSs are used by Battigalli (1996) to define strong

³⁰Consistent systems of conjectures are arrays of probability measures $(\mu_h^i)_{h \in H_i} \in [\Delta(S_{-i})]^{H_i}$ such that, for all $h, \bar{h} \in H_i$, (1) $\mu_h^i(S_{-i}(h)) > 0$, (2) $h \prec \bar{h}$ and $\mu_h^i(S_{-i}(\bar{h})) > 0$ imply $\mu_h^i = \mu_{\bar{h}}^i$. In words, conjectures are consistent with evidence and they are maintained as long as they are not "falsified." Intuitively, it is as

rationalizability.

On the other hand, most works on the epistemic foundations of those same solution concepts use CPSs $\mu^i \in \Delta^{\mathcal{H}_i}(S_{-i})$ with $\mathcal{H}_i = \{S_{-i}(h_i) : h_i \in H_i\}$ as in Section 3, which correspond to standard belief systems; see, e.g., Battigalli & Siniscalchi (1999, 2002), Battigalli & Friedenberg (2012), and Battigalli & De Vito (2021). From a technical point of view, our results establish a bridge between the predictions of solution concepts that were defined with forward consistent belief systems, and the epistemic foundations that have been established for solution procedures defined by means of conditional probability systems.

Finally, our results also connect to known characterizations of rationalizability procedures in sequential games. Shimoji & Watson’s (1998) characterization of strong rationalizability in terms of iterated conditional dominance was proved with forward consistent systems of beliefs. Theorem 5 implies that the characterization holds for the three kinds of belief systems considered here, and similarly for Shimoji’s (2004) result of generic equivalence between justifiability and admissibility.

9 Conclusion

In sequential games, players must update or revise their beliefs about the behavior of co-players upon the arrival of new information regarding their past moves. Various restrictions on how players revise their beliefs have been used in the literature, but the very meaning of these requirements and their implications for optimal planning and strategic reasoning remained somewhat unclear. We put forward and analyze three nested notions of consistency among beliefs at different information sets, and we offer interpretations in terms of players’ introspection.

The weakest notion, forward consistency, only requires belief updating by conditioning, whenever possible, along every path of play. Forward consistency takes the viewpoint of a player who, at each information set, only asks herself how the co-players will behave in the contingencies that she currently deems possible. When surprised by moves that she previously deemed impossible, such a player might form different beliefs depending on her own past moves, not having anticipated how she would have revised her beliefs upon the observation of unexpected moves by the co-players. Our intermediate notion of belief system, standard belief systems, rules this out by requiring that the beliefs at information sets inducing nested information about co-players behavior be related by the chain rule; in particular, having the same information about co-players implies having the same beliefs about them, even if information sets differ because of own past moves. Standard belief systems are isomorphic to

if—in a population-game scenario—each player always holds a Dirac belief about the frequency distribution of strategies in the opponents’ population.

conditional probability systems as in Renyi (1955), defined over the collection of conditioning events concerning co-players' behavior that correspond to information sets. Our strongest notion of consistency, complete consistency, requires instead that beliefs at different information sets can be derived from a conditional probability system over the power set of the space of uncertainty (the set co-players' strategy profiles), which also contains conditioning events that cannot be observed at any point in the game. Such "hypothetical conditioning" can be interpreted as the "mental exercise" of a fully introspective player, who asks herself "What would I believe if, hypothetically, I knew that this event is the case?" and derives her beliefs at information sets accordingly. Complete consistency implies that the relative probabilities of pairs of co-players' strategy profiles are constant across all information sets where such profiles are possible, a property that standard belief systems may violate, as we show by example.

Although complete consistency is much more restrictive than forward consistency for players' belief systems, the main result of the paper shows that justifiability (weak sequential optimality under some belief system) and strong rationalizability are invariant to the assumed notion of consistency of belief systems. The key observation for these results is that forward consistency is sufficient to identify a player's courses of action that are weakly sequentially optimal under a completely consistent belief system.

Appendix A: Omitted proofs of extension results

Proof of Lemma 1

Consider a CPS $\mu \in \Delta^c(\Omega)$ that strongly believes E^n, \dots, E^1 , and fix a *complete CPS* $\tilde{\mu} \in \Delta^*(\Omega)$ that *strongly believes* E^n, \dots, E^1 (such CPS exists, see Remark 2.2 and Lemma A.1 in Battigalli 1996). By the order-extension principle (e.g., Davey & Priestley 2002, p. 32), there is a linear extension of the strict partial order \supset on \mathcal{C} , that is, a strict total order \prec on \mathcal{C} such that $C \supset D \Rightarrow C \prec D$ for all $C, D \in \mathcal{C}$ (Ω is the root of the event tree \mathcal{C} , hence the *minimal* element in this order). Let $<$ denote the corresponding total order on the graph of map $C \mapsto \mu(\cdot|C)$, that is, the set of *indexed* measures induced by μ (think of $\mu(\cdot|C)$ as a *pair* given by conditioning event C and a probability measure on Ω , so that $C \neq D$ implies that $\mu(\cdot|C)$ and $\mu(\cdot|D)$ are different objects even if they correspond to the same measure). With this,

$$\forall C, D \in \mathcal{C}, \quad C \supset D \Rightarrow \mu(\cdot|C) < \mu(\cdot|D).$$

We construct the desired extension $\bar{\mu} \in \Delta^*(\Omega)$ of μ as follows. Let $E^0 := \Omega$, so that $E^n \subseteq \dots \subseteq E^0$ and both μ and the complete CPS $\tilde{\mu}$ strongly believe E^n, \dots, E^0 . For each

$C \neq \emptyset$, let

$$\kappa(C) := \max \{k \in \{0, \dots, n\} : E^k \cap C \neq \emptyset\}$$

denote the highest integer k such that $E^k \cap C \neq \emptyset$ ($\kappa(C)$ is well defined, because $C \cap E^0 = C \neq \emptyset$). Let $M^+(C; \mu)$ denote the set of all indexed measures in the graph of μ that assign strictly positive probability to $E^{\kappa(C)} \cap C$. If $M^+(C; \mu) \neq \emptyset$, derive $\bar{\mu}(\cdot|C)$ by conditioning on C the “earliest” (minimal) indexed measure in $M^+(C; \mu)$ according to $<$ (such measure is unique since $<$ totally orders the graph of μ): if such measure is $\mu(\cdot|\hat{C})$, then $\mu(C \cap \hat{C}|\hat{C}) > 0$ and (recalling that $\mu(\hat{C}|\hat{C}) = 1$)

$$\forall \omega \in C, \quad \bar{\mu}(\omega|C) = \frac{\mu(\omega|\hat{C})}{\mu(C|\hat{C})} = \frac{\mu(\omega|\hat{C})}{\mu(C \cap \hat{C}|\hat{C})}.$$

If $M^+(C; \mu) = \emptyset$, set $\bar{\mu}(\cdot|C) = \tilde{\mu}(\cdot|C)$.

CLAIM 1: $\bar{\mu}$ is a complete CPS. Fix C, D such that $\emptyset \neq D \subset C$ and $\bar{\mu}(D|C) > 0$. Suppose first that $M^+(C; \mu) = \emptyset$. Then $\bar{\mu}(\cdot|C) = \tilde{\mu}(\cdot|C)$. Since $\tilde{\mu}$ strongly believes $E^{\kappa(C)}$, we have $\bar{\mu}(E^{\kappa(C)}|C) = \tilde{\mu}(E^{\kappa(C)}|C) = 1$. Thus, $\bar{\mu}(D|C) > 0$ implies $E^{\kappa(C)} \cap D \neq \emptyset$ and $\kappa(D) = \kappa(C)$ ($\kappa(D) \leq \kappa(C)$ since $D \subset C$). Furthermore, since $D \subset C$, we have $M^+(D; \mu) \subseteq M^+(C; \mu) = \emptyset$ and so $\bar{\mu}(\cdot|D) = \tilde{\mu}(\cdot|D)$ as well. Then the chain rule relates $\bar{\mu}(\cdot|C) = \tilde{\mu}(\cdot|C)$ to $\bar{\mu}(\cdot|D) = \tilde{\mu}(\cdot|D)$, because $\tilde{\mu}$ is a complete CPS. Suppose now that $M^+(C; \mu) \neq \emptyset$. Then $\bar{\mu}(\cdot|C)$ is derived from some $\mu(\cdot|\hat{C}) \in M^+(C; \mu)$; thus, $\mu(E^{\kappa(C)} \cap C|\hat{C}) > 0$ and $E^{\kappa(C)} \cap \hat{C} \neq \emptyset$. Since μ strongly believes $E^{\kappa(C)}$, $\mu(E^{\kappa(C)}|\hat{C}) = 1$. Moreover, $\bar{\mu}(D|C) > 0$ implies $\mu(D|\hat{C}) > 0$. Then, $\mu(E^{\kappa(C)} \cap D|\hat{C}) > 0$ and $E^{\kappa(C)} \cap D \neq \emptyset$, so that $\kappa(D) = \kappa(C)$ and $\mu(\cdot|\hat{C}) \in M^+(D; \mu)$. Furthermore, since $D \subset C$ we have $M^+(D; \mu) \subseteq M^+(C; \mu)$. Since $\mu(\cdot|\hat{C})$ is the minimal indexed measure in $M^+(C; \mu)$, it must also be the minimal one in $M^+(D; \mu)$, and $\bar{\mu}(\cdot|D)$ is derived by conditioning $\mu(\cdot|\hat{C})$ as well. Thus, $\bar{\mu}(\cdot|C)$ and $\bar{\mu}(\cdot|D)$ are derived by conditioning the same measure, which, together with $D \subset C$, implies that they are related by the chain rule.

CLAIM 2: $\bar{\mu}(\cdot|\hat{C}) = \mu(\cdot|\hat{C})$ for all $\hat{C} \in \mathcal{C}$. Fix $\hat{C} \in \mathcal{C}$. Since μ strongly believes $E^{\kappa(\hat{C})}$, we have $\mu(E^{\kappa(\hat{C})}|\hat{C}) = 1 > 0$; hence, $\mu(\cdot|\hat{C}) \in M^+(\hat{C}; \mu) \neq \emptyset$. Thus, $\bar{\mu}(\cdot|\hat{C})$ is derived by conditioning on \hat{C} the minimal indexed measure $\mu(\cdot|\bar{C}) \in M^+(\hat{C}; \mu)$, where $\bar{C} \in \mathcal{C}$ by definition of $M^+(\hat{C}; \mu)$. Moreover, $\hat{C} \subseteq \bar{C}$ (hence, $\bar{C} \preceq \hat{C}$). Indeed, $\mu(\hat{C}|\bar{C}) > 0$ implies $\hat{C} \cap \bar{C} \neq \emptyset$; since \mathcal{C} is an event tree, $\bar{C} \subseteq \hat{C}$ or $\hat{C} \subseteq \bar{C}$. By definition of $M^+(\hat{C}; \mu)$ and by the property of $<$, the linear order defined on the graph of μ , $\hat{C} \subseteq \bar{C}$ must be the case. Then the chain rule for μ implies that $\mu(\cdot|\hat{C})$ is derived by conditioning $\mu(\cdot|\bar{C})$ on \hat{C} , so that $\bar{\mu}(\cdot|\hat{C}) = \mu(\cdot|\hat{C})$.

CLAIM 3: $\bar{\mu}$ strongly believes E^1, \dots, E^n . Fix k and C so that $C \cap E^k \neq \emptyset$. Then $\kappa(C) \geq k$. If $M^+(C; \mu) \neq \emptyset$, then $\bar{\mu}(\cdot|C)$ is derived by conditioning some $\mu(\cdot|\hat{C}) \in M^+(C; \mu)$ on C ; since $\mu(E^{\kappa(C)} \cap C|\hat{C}) > 0$, μ strongly believes $E^{\kappa(C)}$ and $E^{\kappa(C)} \subseteq E^k$, $\bar{\mu}(E^k|C) \geq \bar{\mu}(E^{\kappa(C)}|C) = \frac{\mu(E^{\kappa(C)}|\hat{C})}{\mu(C|\hat{C})} = 1$. If $M^+(C; \mu) = \emptyset$, then $\bar{\mu}(\cdot|C) = \tilde{\mu}(\cdot|C)$; since $\tilde{\mu}$ strongly believes $E^{\kappa(C)}$, $\bar{\mu}(E^k|C) \geq \bar{\mu}(E^{\kappa(C)}|C) = \tilde{\mu}(E^{\kappa(C)}|C) = 1$. ■

9.1 Proof of Lemma 4, Claim 2

First construct an array $(\hat{\mu}^i(\cdot|C))_{C \in \mathcal{H}_i(s_i)}$ as follows. By perfect recall, for each $C \in \mathcal{H}_i(s_i)$ there is a unique \prec_i -minimal $\eta_i(C) \in H_i(s_i)$ such that $S_{-i}(\eta_i(C)) = C$. With this, let $\hat{\mu}^i(\cdot|C) = \mu^i(\cdot|\eta_i(C))$ for every $C \in \mathcal{H}_i(s_i)$. We show that $\hat{\mu}^i$ is a CPS because it satisfies the chain rule. Fix $h_i, \bar{h}_i \in H_i(s_i)$ arbitrarily. If $S_{-i}(h_i) \cap S_{-i}(\bar{h}_i) = \emptyset$, there is nothing to prove. Otherwise, since $\mathcal{H}_i(s_i)$ is an event tree, it is sufficient to consider the case $S(\bar{h}_i) \subseteq S(h_i)$. Since at least one of the alert players is active, $S(\bar{h}_i) = S(h_i)$ implies $\bar{h}_i = h_i$, $\eta_i(S_{-i}(\bar{h}_i)) = \eta_i(S_{-i}(h_i))$ and $\hat{\mu}^i(\cdot|S_{-i}(\bar{h}_i)) = \hat{\mu}^i(\cdot|S_{-i}(h_i))$. Thus, suppose that $S(\bar{h}_i) \subset S(h_i)$. Then $h_i \prec_i \bar{h}_i$ and $\eta_i(S_{-i}(h_i)) \preceq_i h_i \prec_i \eta_i(S_{-i}(\bar{h}_i)) \preceq_i \bar{h}_i$. By construction, $\hat{\mu}(\cdot|S_{-i}(h_i)) = \mu^i(\cdot|\eta_i(S_{-i}(h_i)))$ and $\hat{\mu}(\cdot|S_{-i}(\bar{h}_i)) = \mu^i(\cdot|\eta_i(S_{-i}(\bar{h}_i)))$ with $S_{-i}(\eta_i(S_{-i}(h_i))) = S_{-i}(h_i)$, $S_{-i}(\eta_i(S_{-i}(\bar{h}_i))) = S_{-i}(\bar{h}_i)$. Thus, the forward chain rule for μ^i relates $\hat{\mu}^i(\cdot|S_{-i}(h_i)) = \mu^i(\cdot|\eta_i(S_{-i}(h_i)))$ to $\hat{\mu}^i(\cdot|S_{-i}(\bar{h}_i)) = \mu^i(\cdot|\eta_i(S_{-i}(\bar{h}_i)))$. Since μ^i strongly believes each event $E_{-i}^1, \dots, E_{-i}^n$, it follows from the construction that also $\hat{\mu}^i$ does. ■

Appendix B: Backwards rationalizability

In this Appendix, we show that an invariance result analogous to Theorem 5 holds for backwards rationalizability. Backwards rationalizability characterizes the behavioral implications of rationality and common belief in continuation rationality starting from every information set (cf. Perea 2014, Battigalli & De Vito 2021). In generic games with perfect information, backwards rationalizability coincides with subgame perfect equilibrium and can therefore be computed by backward induction. Catonini & Penta (2022) show that in more general games backwards rationalizability can be computed by means of a convenient Backwards Procedure.³¹ A result by Chen & Micali (2013) on the order-independence of iterated dominance procedures in sequential games implies that the set of strongly rationalizable paths is (weakly) contained in the set of backwards rationalizable paths. Therefore, backwards rationalizability (possibly computed with the backward procedure) can help finding the strongly rationalizable paths.

³¹In a related vein, Perea (2014) introduces the Backward Dominance procedure.

In this section, we focus on games with *observable actions* and consider the definition of backwards rationalizability given by Battigalli & De Vito (2021),³² which is behaviorally equivalent to the one of Perea (2014) for this class of games, as it yields the same *reduced* strategies.³³ To justify this solution concept, Battigalli & De Vito (2021) start from the following epistemic assumptions: Each player is fully introspective about her belief system and plans rationally by folding back, thus obtaining a strategy that is sequentially optimal under her belief system. The player is rational if her behavior complies with her strategy. Players initially believe that co-players are rational. The key assumption is that when a player observes unexpected behavior she keeps believing that the co-players planned rationally and that, even if in the past they deviated from their plans, they will nonetheless comply with their plans in the subgame (according to strong rationalizability, instead, surprised players change their beliefs about the plans that their supposedly rational co-players are implementing). The assumptions concerning higher levels of mutual belief in rationality have a similar flavor.

With this, we want to express the idea that a player in a subgame believes that her co-players are going to behave according to plans having some properties, even if such properties were violated in the past. For example, a nonterminal history $h \in H$ may be unreachable if the co-players' are rational, and yet player i may believe that in the subgame with root h the co-players will implement *continuation strategies* consistent with their rationality *in the subgame*. With this in mind, given a subset of j 's strategies $E_j \subseteq S_j$ and a history $h \in H$, let

$$\chi_j^h(E_j) = \{s_j \in S_j(h) : \exists s'_j \in E_j, \forall \bar{h} \in H, h \preceq \bar{h} \Rightarrow s'_j(\bar{h}) = s_j(\bar{h})\}$$

denote the set of strategies s_j consistent with h that behave like some strategy in E_j in the subgame with root h . As we consider all the co-players, for each profile of strategy subsets $(E_j)_{j \neq i}$, we obtain $\chi_{-i}^h(E_{-i}) = \times_{j \neq i} \chi_j^h(E_j)$, where $E_{-i} = \times_{j \neq i} E_j$. With this, we define the set of belief systems μ^i that “**persistently believe**” in the projection of $E_{-i} = \times_{j \neq i} E_j$ in each subgame:

$$\text{PB}_i(E_{-i}) = \{\mu^i \in \times_{h \in H} \Delta(S_{-i}(h)) : \forall h \in H, \mu^i(\chi_{-i}^h(E_{-i})|h) = 1\}.$$

Note that, unlike the strong belief operator $\text{SB}_i(\cdot)$, operator $\text{PB}_i(\cdot)$ is monotone. Depending on which kind of belief systems we use for the definition of backwards rationalizability, we let Δ_i denote the set of either forward consistent, or standard, or completely consistent belief systems. The main result of this section is that the notions of backwards rationalizability

³²Catonini & Penta (2021) define Backwards Rationalizability for games with incomplete information. Their definition coincides with the definition of Battigalli & De Vito (2021) in games with complete information.

³³Perea (2014) only considers reduced strategies. Battigalli & De Vito (2021) consider full strategies.

obtained by considering these different kinds of belief system are behaviorally equivalent.

Recall that $BR_i^*(\mu^i)$ is the set of strategies of i that are sequentially optimal under belief system μ^i ; given the observability of actions, it is the set of strategies whose continuations maximize i 's expected payoff in each subgame with root $h \in H$ given belief $\mu^i(\cdot|h)$.

Definition 13. Fix $\Delta = (\Delta_i)_{i \in I} \in \left\{ (\Delta_F^{H_i}(S_{-i}))_{i \in I}, (\Delta^{H_i}(S_{-i}))_{i \in I}, (\Delta_C^{H_i}(S_{-i}))_{i \in I} \right\}$. For each $i \in I$, let

$$\hat{S}_i^{\Delta,1} = \{s_i \in S_i : \exists \mu^i \in \Delta_i, s_i \in BR_i^*(\mu^i)\},$$

and define recursively, for each $m \in \mathbb{N}$ and $i \in I$,

$$\hat{S}_i^{\Delta,m+1} = \{s_i \in S_i : \exists \mu^i \in \Delta_i \cap \text{PB}_i(\hat{S}_{-i}^{\Delta,m}), s_i \in BR_i^*(\mu^i)\}.$$

The set of **backwards rationalizable** strategies for i **with forward consistent** (respectively, **standard, completely consistent**) **belief systems** is $\hat{S}_i^{\Delta,\infty} = \bigcap_{k=1}^{\infty} \hat{S}_i^{\Delta,k}$ where $\Delta = (\Delta_F^{H_i}(S_{-i}))_{i \in I}$ (respectively, $\Delta = (\Delta^{H_i}(S_{-i}))_{i \in I}$, $\Delta = (\Delta_C^{H_i}(S_{-i}))_{i \in I}$).

Now we show the behavioral equivalence result for backwards rationalizability.

Theorem 8. *The backwards rationalizability procedures of Definition 13 coincide.*

To prove this result, it is first convenient to introduce a stronger notion of “**persistent belief**” in E_{-i} , that is, the belief at each h that the co-players will follow, or would have followed some strategy profile in E_{-i} starting from every $h' \in H$ that does not precede h , including the histories h' that are counterfactual given h . Let

$$\hat{\chi}_j^h(E_j) = \{s_j \in S_j(h) : \exists \bar{s}_j \in E_j, \forall h' \in H, h' \not\prec h \Rightarrow s_j(h') = \bar{s}_j(h')\}$$

and

$$\widehat{\text{PB}}_i(E_{-i}) = \{\mu^i \in \times_{h \in H} \Delta(S_{-i}(h)) : \forall h \in H, \mu^i(\hat{\chi}_{-i}^h(E_{-i})|h) = 1\},$$

where $E_{-i} = \times_{j \neq i} E_j$ and $\hat{\chi}_{-i}^h(E_{-i}) = \times_{j \neq i} \hat{\chi}_j^h(E_j)$.

We proceed by establishing some preliminary lemmas.

Lemma 6. *Fix a strategy \bar{s}_i , a subset $E_{-i} \subseteq S_{-i}$, and a forward consistent belief system $\bar{\mu}^i$ in $\widehat{\text{PB}}_i(E_{-i})$. Then there exists a completely consistent μ^i such that (i) $\mu^i(\cdot|S_{-i}(h)) = \bar{\mu}^i(\cdot|h)$ for all $h \in H_i(\bar{s}_i)$ and (ii) $\mu^i \in \text{PB}_i(E_{-i})$.*

Proof: To economize on constructions, we take the same complete CPS μ^i as the one of Lemma 4, but assume that $\bar{\mu}^i$ is in $\widehat{\text{PB}}_i(E_{-i})$ instead of $\text{SB}_i(E_{-i})$.

CLAIM: $\mu^i \in \text{PB}_i(E_{-i})$. Fix any $h \in H$. If $\mu^i(\cdot|S_{-i}(h)) = \tilde{\mu}^i(\cdot|S_{-i}(h))$ or $\mu^i(\cdot|S_{-i}(h))$ was derived from ν , we have $\mu^i(\chi_{-i}^h(E_{-i})|S_{-i}(h)) = 1$. Suppose now that $\mu^i(\cdot|S_{-i}(h))$ is derived from some $\bar{\mu}^i(\cdot|\bar{h})$ in the ordering; it cannot be that $h \prec \bar{h}$, so we have $h \not\prec \bar{h}$. Now, note that, since $\mu^i(\cdot|S_{-i}(h))$ is derived from $\bar{\mu}^i(\cdot|\bar{h})$ and $\bar{\mu}^i \in \text{PB}_i(E_{-i})$, we have $\text{Supp } \mu^i(\cdot|S_{-i}(h)) \subseteq \hat{\chi}_{-i}^{\bar{h}}(E_{-i})$. Moreover, we claim that $h \not\prec \bar{h}$ implies $\hat{\chi}_{-i}^{\bar{h}}(E_{-i}) \subseteq \chi_{-i}^h(E_{-i})$. Suppose $h \not\prec \bar{h}$ and fix $s_{-i}^* \in \hat{\chi}_{-i}^{\bar{h}}(E_{-i})$. Then there exists $s_{-i} \in E_{-i}$ such that $s_{-i}(h') = s_{-i}^*(h')$ for all $h' \not\prec \bar{h}$. Since $h \not\prec \bar{h}$, every $h' \succeq h$ satisfies $h' \not\prec \bar{h}$, and so, $s_{-i}(h') = s_{-i}^*(h')$ for all $h' \succeq h$. Hence we have $s_{-i}^*|h = s_{-i}|h$ for some $s_{-i} \in E_{-i}$, that is to say, $s_{-i}^* \in \chi_{-i}^h(E_{-i})$. So, $\text{Supp } \mu^i(\cdot|S_{-i}(h)) \subseteq \chi_{-i}^h(E_{-i})$ for all $h \in H$, i.e., $\mu^i \in \text{PB}_i(E_{-i})$. ■

Given two belief systems $\mu^i, \bar{\mu}^i$ and a nonterminal history $h \in H$, we write $\mu^i(\cdot|h) \simeq \bar{\mu}^i(\cdot|h)$ if $\mu^i(S_{-i}(z)|h) = \bar{\mu}^i(S_{-i}(z)|h)$ for every $z \succ h$. We write $\mu^i \simeq \bar{\mu}^i$ if $\mu^i(\cdot|h) \simeq \bar{\mu}^i(\cdot|h)$ for all $h \in H$.

Lemma 7. *Fix a subset $E_{-i} \subseteq S_{-i}$, a forward consistent belief system $\bar{\mu}^i \in \text{PB}_i(E_{-i})$, and a strategy \bar{s}_i . There exists a completely consistent $\mu^i \in \text{PB}_i(E_{-i})$ such that $\mu^i(\cdot|h) \simeq \bar{\mu}^i(\cdot|h)$ for every $h \in H_i(E_i)$.*

Proof: Fix a forward consistent belief system $\bar{\mu}^i \in \text{PB}_i(E_{-i})$. Note that, for each h , $\hat{\chi}_{-i}^h(\bar{S}_{-i}) \subseteq \chi_{-i}^h(E_{-i})$. With this, for each $h \in H$, define a map $\zeta_{-i}^h : \chi_{-i}^h(E_{-i}) \rightarrow \hat{\chi}_{-i}^h(E_{-i})$, which is the identity on $\hat{\chi}_{-i}^h(E_{-i})$ and is such that for each $s_{-i} \in \chi_{-i}^h(E_{-i})$, $\zeta_{-i}^h(s_{-i})|h = s_{-i}|h$. Note that for all $h' \succ h$ and $s_{-i} \in \chi_{-i}^h(E_{-i})$, $s_{-i} \in S_{-i}(h')$ if and only if $\zeta_{-i}^h(s_{-i}) \in S_{-i}(h')$. It is convenient to impose the following consistency property: for all $h' \succ h$ and $s_{-i} \in S_{-i}(h') \subseteq S_{-i}(h)$, set $\zeta_{-i}^{h'}(s_{-i}) = \zeta_{-i}^h(s_{-i})$.

Take the belief system $(\tilde{\mu}^i(\cdot|h))_{h \in H}$ defined as follows: For every $h \in H$ and $\hat{s}_{-i} \in \hat{\chi}_{-i}^h(\bar{S}_{-i})$, let

$$\tilde{\mu}^i(\hat{s}_{-i}|h) = \bar{\mu}^i(\{s_{-i} : \zeta_{-i}^h(s_{-i}) = \hat{s}_{-i}\}|h) = \bar{\mu}^i(\zeta_{-i}^{h,-1}(\hat{s}_{-i})|h).$$

Notice that for every $h \in H$, $\tilde{\mu}^i(S_{-i}(h')|h) = \bar{\mu}^i(S_{-i}(h')|h)$ for all $h' \succ h$ ($h' \in \bar{H}$), which implies $\tilde{\mu}^i \simeq \bar{\mu}^i$. Also, since $\bar{\mu}^i \in \text{PB}_i(\bar{S}_{-i})$, we have

$$\tilde{\mu}^i(\hat{\chi}_{-i}^h(E_{-i})|h) = \bar{\mu}^i(\zeta_{-i}^{h,-1}(\hat{\chi}_{-i}^h(E_{-i}))|h) = \bar{\mu}^i(\chi_{-i}^h(E_{-i})|h) = 1$$

for all $h \in H$. Thus, $\tilde{\mu}^i \in \text{PB}_i(E_{-i})$. We show that $\tilde{\mu}^i$ is a forward consistent belief system. Fix $h \prec h'$ and $E_{-i} \subseteq S_{-i}(h')$ such that $\tilde{\mu}^i(E_{-i}|h) > 0$. Then, note first that $\zeta_{-i}^{h,-1}(E_{-i}) \subseteq S_{-i}(h')$ and that $\bar{\mu}^i(\zeta_{-i}^{h,-1}(E_{-i})|h) = \tilde{\mu}^i(E_{-i}|h) > 0$. Thus, the forward chain

rule of $\bar{\mu}^i$ applies and implies:

$$\tilde{\mu}^i(E_{-i}|h') = \bar{\mu}^i(\varsigma_{-i}^{h',-1}(E_{-i})|h') = \bar{\mu}^i(\varsigma_{-i}^{h,-1}(E_{-i})|h') = \frac{\bar{\mu}^i(\varsigma_{-i}^{h,-1}(E_{-i})|h)}{\bar{\mu}^i(S_{-i}(h')|h)} = \frac{\tilde{\mu}^i(E_{-i}|h)}{\tilde{\mu}^i(S_{-i}(h')|h)},$$

where the second equality holds by the consistency condition of the maps $(\varsigma_{-i}^h)_{h \in H}$; the third equality by the forward chain rule of $\bar{\mu}^i$; and the fourth equality by definition of $\tilde{\mu}^i$ and $\tilde{\mu}^i(S_{-i}(h')|h) = \bar{\mu}^i(S_{-i}(h')|h)$.

To summarize, we have a forward consistent belief system $\tilde{\mu}^i \in \text{PB}_i(E_{-i})$ such that $\tilde{\mu}^i(\cdot|h) \simeq \bar{\mu}^i(\cdot|h)$ for every $h \in H$. Then, by Lemma 6, there exists a completely consistent μ^i in $\text{PB}_i(E_{-i})$ such that $\mu^i(\cdot|h) = \tilde{\mu}^i(\cdot|h)$ for all $h \in H_i(\bar{s}_i)$. Hence, $\mu^i(\cdot|h) \simeq \bar{\mu}^i(\cdot|h)$ for all $h \in H_i(\bar{s}_i)$. ■

Proof of Theorem 8: Denote by $(\tilde{S}^n)_{n \in \mathbb{N}_0}$ and $(\hat{S}^n)_{n \in \mathbb{N}_0}$ the backwards rationalizability algorithms with, respectively, forward consistent and completely consistent belief systems.

We show the following inductive hypothesis: for each i , (i) $\hat{S}_i^n \subseteq \tilde{S}_i^n$, and (ii) for all $h \in H$ there is a map $\varsigma_i^{n,h} : \chi_i^h(\tilde{S}_i^n) \rightarrow \chi_i^h(\hat{S}_i^n)$ such that, for each $s_i \in \chi_i^h(\tilde{S}_i^n)$, $\varsigma_i^{n,h}(s_i)(h') = s_i(h')$ for all $h' \in H_i(s_i)$. The basis step is trivial, we show the statement holds at step $n + 1$.

Fix any $i \in I$. First, note that for any consistent CPS μ^i in $\text{PB}_i(\hat{S}_{-i}^n)$, by (i) of the inductive hypothesis and monotonicity of the $\text{PB}_i(\cdot)$ operator, we have $\mu^i \in \text{PB}_i(\tilde{S}_{-i}^n)$, and since a consistent CPS is a forward consistent belief system, we have (i) $\hat{S}_i^{n+1} \subseteq \tilde{S}_i^{n+1}$.

Now, take any $\bar{h} \in H$ and $s_i \in \chi_i^{\bar{h}}(\tilde{S}_i^{n+1})$. Then by definition there is a forward consistent belief system $\mu^i \in \text{PB}_i(\tilde{S}_{-i}^n)$ such that $s'_i|\bar{h} = s_i|\bar{h}$ for some $s'_i \in BR_i^*(\mu^i)$.

Define a belief system $(\hat{\mu}^i(\cdot|h))_{h \in H} \in \times_{h \in H} \Delta(S_{-i}(h))$ as follows: For $h = \emptyset$ and all $h \neq \emptyset$ such that $\mu^i(S_{-i}(h)|\bar{h}) = 0$ for all $\bar{h} \prec h$, set:

$$\forall \hat{s}_{-i} \in \chi_{-i}^h(\hat{S}_{-i}^n), \quad \hat{\mu}^i(\hat{s}_{-i}|h) = \mu^i(\varsigma_{-i}^{n,h,-1}(\hat{s}_{-i})|h).$$

We have the following properties: since for every $h' \succ h$ and $s_{-i} \in \chi_{-i}^h(\tilde{S}_{-i}^n)$, $s_{-i} \in S_{-i}(h')$ if and only if $\varsigma_{-i}^{n,h}(s_{-i}) \in S_{-i}(h')$, we have (a) $\hat{\mu}^i(S_{-i}(h')|h) = 0$ if and only if $\mu^i(S_{-i}(h')|h) = 0$ and (b) $\mu^i \simeq \hat{\mu}^i$. Moreover, since $\mu^i \in \text{PB}_i(\tilde{S}_{-i}^n)$,

$$\hat{\mu}^i(\chi_{-i}^h(\hat{S}_{-i}^n)|h) = \mu^i(\varsigma_{-i}^{n,h,-1}(\chi_{-i}^h(\hat{S}_{-i}^n))|h) = \mu^i(\chi_{-i}^h(\tilde{S}_{-i}^n)|h) = 1.$$

Thus, deriving all other probability measures by forward conditioning, we indeed obtain a forward consistent belief system $\hat{\mu}^i$. It can be checked using a forward conditioning argument that $\hat{\mu}^i \simeq \mu^i$, so that $BR_i^*(\mu^i) = BR_i^*(\hat{\mu}^i)$. Moreover, $\hat{\mu}^i$ belongs to $\text{PB}_i(\hat{S}_{-i}^n)$. To see this, take any $h' \succ h$. This implies $\chi_{-i}^h(\hat{S}_{-i}^n) \cap S_{-i}(h') \subseteq \chi_{-i}^{h'}(\hat{S}_{-i}^n)$,³⁴ and so, if $\hat{\mu}^i(\cdot|h')$ is derived

³⁴See Lemma 4.(i) in Appendix A of Battigalli & De Vito (2021).

by conditioning from $\hat{\mu}^i(\cdot|h)$, because $\text{Supp } \hat{\mu}^i(\cdot|h) \subseteq \chi_{-i}^h(\hat{S}_{-i}^n)$, we have:

$$\text{Supp } \hat{\mu}^i(\cdot|h') \subseteq \chi_{-i}^h(\hat{S}_{-i}^n) \cap S_{-i}(h') \subseteq \chi_{-i}^{h'}(\hat{S}_{-i}^n).$$

To summarize, we have a forward consistent belief system $\hat{\mu}^i$ such that $\hat{\mu}^i \in \text{PB}_i(\hat{S}_{-i}^n)$ and $\hat{\mu}^i \simeq \mu^i$. Then by Lemma 7 there exists a consistent CPS $\tilde{\mu}^i \in \text{PB}_i(\hat{S}_{-i}^n)$ with $\tilde{\mu}^i(\cdot|h) \simeq \mu^i(\cdot|h)$ for all $h \in H_i(s'_i)$. This in turn implies that there is a strategy $\tilde{s}_i \in BR_i^*(\tilde{\mu}^i)$ such that, for each $h \in H_i(s'_i)$ with $h \succeq \bar{h}$, $\tilde{s}_i(h) = s'_i(h) = s_i(h)$. Thus, $\tilde{s}_i \in \hat{S}_i^{n+1}$ and any $\bar{s}_i \in S_i(\bar{h})$ with $\bar{s}_i|\bar{h} = \tilde{s}_i|\bar{h}$ belongs to $\chi_i^h(\hat{S}_i^{n+1})$, so that we can let $\zeta_i^{n+1, \bar{h}}(s_i) = \bar{s}_i$. ■

Appendix C: Comparison with Kreps & Wilson

9.2 Complete consistency and sequences

We let $\Delta^\circ(\Omega) \subseteq \Delta(\Omega)$ denote the set of strictly positive probability measures on Ω .

For any complete CPS $\mu \in \Delta^*(\Omega)$, say that x is μ -infinitely more likely than y , written $x \ggg^\mu y$, if $\mu(x|\{x, y\}) = 1$ (i.e., $\mu(y|\{x, y\}) = 0$). The chain rule implies that the (typically incomplete) binary relation \ggg^μ is transitive. Therefore, μ yields an ordered partition $(\Omega_0^\mu, \dots, \Omega_{L(\mu)}^\mu)$, where—for $0 \leq k < \ell \leq L(\mu)$ —each element of Ω_k^μ is μ -infinitely more likely of each element of Ω_ℓ^μ . With this, $\mu(\cdot|\Omega_0^\mu) = \mu(\cdot|\Omega)$ is the unconditional belief. More generally, let $k(\mu, C)$ denote the lowest index in $\{0, \dots, L(\mu)\}$ such that $C \cap \Omega_k \neq \emptyset$; then $\mu(C \cap \Omega_{k(\mu, C)}|C) = 1$ and

$$\mu(\cdot|C \cap \Omega_{k(\mu, C)}) = \mu(\cdot|C),$$

which implies the following:

Remark 11. A complete CPS $\mu \in \Delta^*(\Omega)$ is determined by the binary conditional probability values $(\mu(\cdot|\{x, y\}))_{x, y \in \Omega}$. Indeed, if two complete CPSs μ and ν are such that

$$(\mu(\cdot|\{x, y\}))_{x, y \in \Omega} = (\nu(\cdot|\{x, y\}))_{x, y \in \Omega}$$

then $(\Omega_0^\mu, \dots, \Omega_{L(\mu)}^\mu) = (\Omega_0^\nu, \dots, \Omega_{L(\nu)}^\nu)$, which implies $\mu = \nu$.

Ordered partitions can also be used to relate complete CPSs to converging sequences of strictly positive probability measures. First note that if a sequence $(\nu_n)_{n \in \mathbb{N}} \in \Delta^\circ(\Omega)^\mathbb{N}$ converges and yields an array $(\mu(\cdot|C))_{\emptyset \neq C \subseteq \Omega} \in \times_{\emptyset \neq C \subseteq \Omega} \Delta(C)$ in the limit, that is,

$$\forall C, \forall \omega \in C, \mu(\omega|C) = \lim_{n \rightarrow \infty} \frac{\nu_n(\omega)}{\nu_n(C)}, \quad (1)$$

then $(\mu(\cdot|C))_{\emptyset \neq C \subseteq \Omega}$ must be a complete CPS. Next note that, for every complete CPS μ , the sequence

$$(\nu_n)_{n \in \mathbb{N}} = \left(K_n \sum_{k=0}^{L(\mu)} n^{-k} \mu(\cdot | \Omega_k^\mu) \right)_{n \in \mathbb{N}} \in \Delta^\circ(\Omega)^\mathbb{N}$$

(where $K_n = \left(\sum_{k=0}^{L(\mu)} n^{-k} \right)^{-1}$ is a normalizing constant) converges and yields μ in the limit. Therefore:

Remark 12. An array of probability measures $\mu = (\mu(\cdot|C))_{\emptyset \neq C \subseteq \Omega} \in \times_{\emptyset \neq C \subseteq \Omega} \Delta(C)$ is a complete CPS if and only if there is a convergent sequence of strictly positive probability measures $(\nu_n)_{n \in \mathbb{N}}$ that yields μ in the limit as in eq. (1).

9.3 Kreps & Wilson's systems of beliefs

An *assessment* à la Kreps & Wilson (1982) is a pair (ξ, β) , where $\xi \in \times_{i \in I} (\times_{h_i \in H_i} \Delta(h_i))$ is a system of beliefs over non-terminal nodes conditional the information sets containing them, and β is a profile of behavior (i.e., locally randomized) strategies. Kreps & Wilson's notion of consistency (used to define sequential equilibrium, a refinement of the Nash equilibrium concept) requires that the beliefs of different players “agree” and that they satisfy what amounts to a strong across-players independence property. Such restrictions are not germane to our (correlated) rationalizability analysis. Thus, we make them mute by considering **personal** assessments in **two-player** games.

A *personal assessment* of player i in a two-player game is a pair

$$(\xi^i, \beta_{-i}) \in (\times_{h_i \in H_i} \Delta(h_i)) \times (\times_{h_{-i} \in H_{-i}} \Delta(\mathcal{A}_{-i}(h_{-i}))),$$

where ξ^i is a system of beliefs over nodes/histories in the information sets of i and β_{-i} is a behavior strategy of the co-player $-i$, which we may interpret as a probabilistic conjecture of i about $-i$. A personal assessment (ξ^i, β_{-i}) is consistent in the sense of Kreps & Wilson if there exists a converging sequence of strictly positive behavior strategies $(\beta_{-i,n})_{n \in \mathbb{N}} \in (\times_{h_{-i} \in H_{-i}} \Delta^\circ(\mathcal{A}_{-i}(h_{-i})))^\mathbb{N}$ such that $\beta_{-i} = \lim_{n \rightarrow \infty} \beta_{-i,n}$ and

$$\xi^i(x|h_i) = \lim_{n \rightarrow \infty} \frac{\mathbb{P}_{\beta_i^\circ, \beta_{-i,n}}(x)}{\mathbb{P}_{\beta_i^\circ, \beta_{-i,n}}(h_i)},$$

for all $h_i \in H_i$ and $x \in h_i$, where \mathbb{P}_β denotes the probability of histories induced by randomized strategy profile (pair) β , and β_i° is an arbitrary strictly positive randomized strategy of i .³⁵ By Kuhn Theorem (Kuhn 1953, Theorem 4), the latter condition is equivalent to the

³⁵By perfect recall, it does not matter which one, because the factors involving β_i° in the numerator and

existence of a converging sequence of strictly positive mixed strategies $(\nu_{-i,n})_{n \in \mathbb{N}} \in \Delta^\circ(S_{-i})^{\mathbb{N}}$ such that $\nu_{\beta_{-i}} = \lim_{n \rightarrow \infty} \nu_{-i,n}$ and

$$\xi^i(x|h_i) = \lim_{n \rightarrow \infty} \frac{\mathbb{P}_{\nu_i^\circ, \nu_{-i,n}}(x)}{\mathbb{P}_{\nu_i^\circ, \nu_{-i,n}}(h_i)},$$

where $\nu_{\beta_{-i}}$ is the mixed strategy induced by β_{-i} according to Kuhn's transformation and ν_i° is an arbitrary strictly positive randomized strategy.³⁶

A personal assessment (ξ^i, β_{-i}) yields a conditional probability measure $\mathbb{P}_{\xi^i, \beta_{-i}, \beta_i}(\cdot|h_i)$ over paths for every behavior strategy β_i of i conditional on every information set h_i of i : if $z = (x, a^{\ell(x)+1}, \dots, a^{\ell(z)})$ with $x \in h_i$ of length $\ell(x)$,

$$\mathbb{P}_{\xi^i, \beta_{-i}, \beta_i}(z|h_i) = \left(\prod_{k=1}^{\ell(z)-\ell(x)} \prod_{j \in \iota(x)} \beta_j(a_j^{\ell(x)+k} | h_j^{\ell(x)+k}) \right) \xi^i(x|h_i).$$

Similarly, a system of beliefs μ^i about the co-player's strategies yields a conditional probability measure $\mathbb{P}_{\mu^i, s_i}(\cdot|h_i)$ over paths for every strategy s_i of i conditional on every information set h_i of i : let $Z(h_i)$ denote the set of terminal nodes/paths following some node in h_i , then

$$\mathbb{P}_{\mu^i, s_i}(z|h_i) = \begin{cases} 0 & \text{if } s_i \notin S_i(z) \text{ or } z \notin Z(h_i), \\ \mu^i(S_{-i}(z)) & \text{if } s_i \in S_i(z) \text{ and } z \in Z(h_i). \end{cases}$$

Note that pure strategies are special cases of degenerate behavior strategies. With a slight abuse of notation we use the same symbol s_i for a pure strategy of i and the corresponding degenerate behavior strategy of i . With this, by Remark 6 we obtain:

Remark 13. A belief system μ^i about the strategies of $-i$ is completely consistent as per Definition 6 if and only if there exists a consistent personal assessment (ξ^i, β_{-i}) such that

$$\mathbb{P}_{\mu^i, s_i}(z|h_i) = \mathbb{P}_{\xi^i, \beta_{-i}, s_i}(z|h_i)$$

for all $h_i \in H_i$, $z \in Z(h_i)$, and $s_i \in S_{-i}(h_i)$.

References

- [1] BATTIGALLI, P. 1996. "Strategic Rationality Orderings and the Best Rationalization Principle," *Games and Economic Behavior* 13: 178-200.

denominator cancel out.

³⁶As above, by perfect recall, it does not matter which one, because the factors involving ν_i° in the numerator and denominator cancel out.

- [2] BATTIGALLI, P. 1997. "On Rationalizability in Extensive Games," *Journal of Economic Theory* 74 (1): 40-61.
- [3] BATTIGALLI, P. AND DE VITO, N. 2021. "Beliefs, Plans and Perceived Intentions in Games," *Journal of Economic Theory* 195: 105283.
- [4] BATTIGALLI, P. AND DUFWENBERG, M. 2022. "Belief-Dependent Motivations and Psychological Game Theory" (with M. Dufwenberg), *Journal of Economic Literature*, in print.
- [5] BATTIGALLI P. AND FRIEDENBERG A. 2012. "Forward Induction Reasoning Revisited," *Theoretical Economics*, 7, 57-98.
- [6] BATTIGALLI, P. AND GENEROSO, N. 2021. "Information Flows and Memory in Games," IGER w.p. 678, Bocconi University.
- [7] BATTIGALLI, P. AND SINISCALCHI, M. 1999. "Hierarchies of Conditional Beliefs and Interactive Epistemology in Dynamic Games," *Journal of Economic Theory* 88: 188-230.
- [8] BATTIGALLI, P. AND SINISCALCHI, M. 2002. "Strong Belief and Forward Induction Reasoning," *Journal of Economic Theory* 106: 356-391.
- [9] BATTIGALLI, P. AND SINISCALCHI, M. 2003. "Rationalization and Incomplete Information," *BE Journals in Theoretical Economics (Advances)* 1:Article 3.
- [10] BEN-PORATH, E. 1997. "Rationality, Nash Equilibrium and Backwards Induction in Perfect-Information Games," *The Review of Economic Studies* 64: 23-46.
- [11] CATONINI, E., 2021. "Self-enforcing agreements and forward induction reasoning," *Review of Economic Studies*, 88, 610-642.
- [12] CATONINI, E., 2022. "A Dutch-Book Argument for Belief Consistency," arXiv:2202.1012 [econ.TH].
- [13] CATONINI, E. AND PENTA A. 2022. "Backward Induction Reasoning beyond Backward Induction," Working Paper.
- [14] DAVEY, B.A., AND PRIESTLEY H.A. (2002): *Introduction to Lattices and Order*. Cambridge UK: CUP (2nd Edition).
- [15] KREPS, D. AND WILSON, R. 1982. "Sequential Equilibria," *Econometrica* 50 (4): 863-894.

- [16] KUHN, H. 1953. "Extensive Games and the Problem of Information," in *Contributions to the Theory of Games*, Vol. 2. Princeton: Princeton University Press, 193-216.
- [17] MYERSON, R. 1986. "Multistage Games with Communication," *Econometrica* 54 (2): 323-358.
- [18] OSBORNE, M., AND RUBINSTEIN, A. 1994. *A Course in Game Theory*, Cambridge (MA): MIT Press.
- [19] PEARCE, D. 1984. "Rationalizable Strategic Behavior and the Problem of Perfection," *Econometrica* 52 (4): 1029-1050.
- [20] PENTA, A. 2015. "Robust Dynamic Implementation," *Journal of Economic Theory* 160: 280-316.
- [21] PEREA, A. 2002. "A Note on the One-Deviation Property in Extensive Form Games," *Games and Economic Behavior* 40: 322-338.
- [22] PEREA, A. 2014. "Belief in the Opponents' Future Rationality," *Games and Economic Behavior* 83: 231-254.
- [23] RENYI, A. 1955. "On a New Axiomatic Theory of Probability," *Acta Mathematica Academiae Scientiarum Hungaricae* 6: 285-335.
- [24] RUBINSTEIN, A. 1991. "Comments on the Interpretation of Game Theory," *Econometrica* 48: 909-924.
- [25] SELTEN, R. 1975. "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games," *International Journal of Game Theory* 4: 25-55.
- [26] SHIMOJI, M. AND WATSON, J. 1998. "Conditional Dominance, Rationalizability, and Game Forms," *Journal of Economic Theory* 83: 161-195.
- [27] SHIMOJI, M. 2004. "On the Equivalence of Weak Dominance and Sequential Best Response," *Games and Economic Behavior* 48: 385-402.
- [28] SINISCALCHI, M. 2022. "Structural Rationality in Dynamic Games," *Econometrica*, forthcoming.
- [29] VON NEUMANN, J. AND MORGENSTERN, O. 1944. *Theory of Games and Economic Behavior*, Princeton: Princeton University Press.