Rational play in games: A behavioral approach

Giacomo Bonanno

Department of Economics, University of California, Davis, USA
gfbonanno@ucdavis.edu

Abstract

We argue in favor of a departure from the standard equilibrium approach in game theory in favor of the less ambitious goal of describing only the *actual behavior* of rational players. We investigate the notion of rationality in behavioral models of extensive-form games (allowing for imperfect information), where a state is described in terms of a play of the game instead of a strategy profile. The players' beliefs are specified only at reached decision histories and are modeled as pre-choice beliefs, allowing us to carry out the analysis without the need for (objective or subjective) counterfactuals. The analysis is close in spirit to the literature on self-confirming equilibrium, but it does not rely on the notion of strategy. We also provide a characterization of rational play that is compatible with pure-strategy Nash equilibrium.

October 29, 2021

1 Introduction

What constitutes a rational solution of an extensive-form game? Different approaches have been put forward in the literature.

1. The equilibrium approach. Within this approach the notion of rationality is captured through the notion of Nash equilibrium or one of its refinements.

Consider, for example, the game of Figure 1 and the Nash equilibrium (a, a, M).

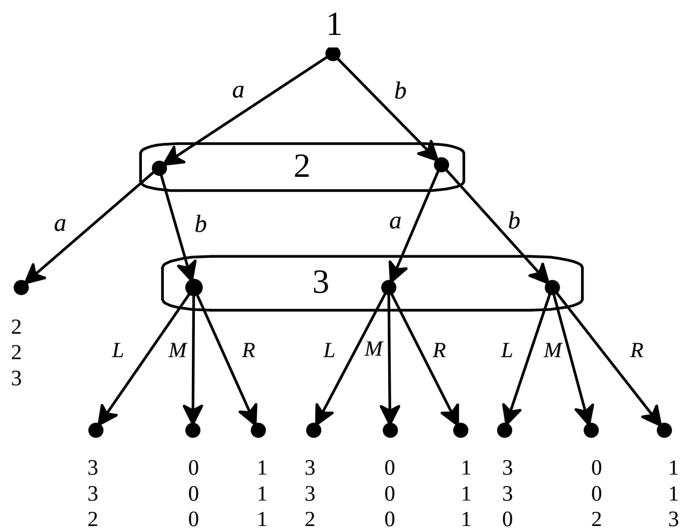

Figure 1: An extensive game with imperfect information.

The strategy of each player is a best response to the strategies of her opponents and can thus be taken to be a rational choice. However, how should Player 3's strategy M be interpreted? Just like any other strategy of Player 3, M is a rational "choice" merely because Player 3's information set is not reached (given that Players 1 and 2 play (a, a)). If M is to be interpreted as the counterfactual "Player 3 would play M if her information set were to be reached" then one can question the rationality of M on the basis of the fact that R is strictly better than M at each history in Player 3's information set. Thus refinements of Nash equilibrium, such as sequential equilibrium (Kreps and Wilson (1982)) and perfect Bayesian equilibrium (Fudenberg and Tirole (1991), Bonanno (2013b)), define a solution not merely as a strategy profile but as an assessment $\langle \sigma, \mu \rangle$, where σ is a strategy profile and μ is a "system of beliefs" consisting of a list of probability distributions, one for every information set, representing the (possibly counterfactual) beliefs of the corresponding players at those information sets. The notion of sequential rationality then restricts choices at information sets to those that are optimal given the specified beliefs and the strategies of the opponents. Sequential rationality rules out strategy M for Player 3 in the game of Figure 1.

Consider now a different Nash equilibrium, namely (a, a, R) together with a system of beliefs μ that assigns probability 1 to histories a and bb. Such an assessment satisfies sequential rationality and constitutes a rational solution if one subscribes to the notion of weak sequential equilibrium (Myerson 1991, p. 170) or perfect Bayesian equilibrium (as defined in Bonanno (2013b)) but not if one subscribes to the notion of sequential equilibrium. The latter requires beliefs to satisfy a property called "consistency" which entails that $\mu(bb) = 0$.[1] But if Player 3 assigns zero probability to history bb then the only sequentially rational choice is L. Thus there is no sequential equilibrium where the strategy profile is (a, a, R). In my opinion, there are no compelling reasons to rule out the belief $\mu(bb) = 1$ and thus (a, a, R) can indeed be considered a rational solution of the game of Figure 1, for the following reason. For strategic-form games, Aumann (Aumann 1987, p. 16) put forward the thesis – which is by now generally accepted – that Player 3, while recognizing that Players 1 and 2 act independently, may nevertheless have beliefs about their choices that display correlation, because she believes that some unobserved common factor has helped determine the choices of both of her opponents. But then a similar argument can be applied to the game of Figure 1: having observed that there was a deviation from the play aa, Player 3 can reasonably come to believe that some unobserved common factor has led both of her opponents to deviate from a to b, that is, she can reasonably form the belief $\mu(bb) = 1$. This argument calls into question restrictions on beliefs that are incorporated in several refinements of Nash equilibrium.

2. The self-confirming equilibrium approach. A second strand of the literature identifies rational play in extensive-form games with the notion of self-confirming equilibrium, introduced in Fudenberg and Levine (1993).[2] A self-confirming equilibrium is a strategy profile satisfying the property that

[1]As originally defined by Kreps and Wilson (Kreps and Wilson (1982)), consistency is a topological notion, but it can be redefined in a way that does not make any reference to sequences of completely mixed strategies: see Bonanno (2016).

[2]Similar notions are put forward in Battigalli and Guaitoli (1997), Greenberg (2000), Greenberg et al. (2009). Dekel et al. (1999; 2002) provide a refinement of self-confirming equilibrium that imposes constraints on the players' beliefs about what actions an opponent could take at an off-path information set and Fudenberg and Kamada (2015) provide a generalization of the notion of self-confirming equilibrium. The related expression 'conjectural equilibrium' is mostly used in the context of strategic-form games; it was introduced in this context by Battigalli (1987), Gilli (1987) and is defined as a situation where each player's strategy is a best response to a conjecture about the

each player's strategy is a best response to her beliefs about the strategies of her opponents, and each player's beliefs are correct along the equilibrium play. At a self-confirming equilibrium no player receives information that contradicts her beliefs, even though her beliefs about play at off-path information sets may be incorrect. For example, in the game of Figure 1, the strategy profile (a, a, R) is a self-confirming equilibrium if Players 1 and 2 share the belief that Player 3 would play R if her information set were to be reached. Note, however, that at a self-confirming equilibrium different players are allowed to have different beliefs about the hypothetical choice of a third player at an unreached information set. As a consequence, a self-confirming equilibrium need not be a Nash equilibrium. Consider, for example, the game of Figure 2, taken from (Fudenberg and Levine 1993, p. 533).

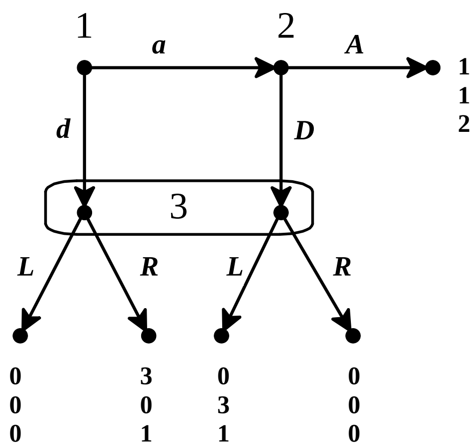

Figure 2: The play aA is consistent with the notion of self-confirming equilibrium.

Both (a, A, L) and (a, A, R) are self-confirming equilibria sustained by Player 1's belief (correct in the former, erroneous in the latter) that Player 3 would play L and Player 2's belief (erroneous in the former, correct in the latter) that Player 3

other players' strategies and any information acquired after the play of the game does not induce the player to change her conjecture. It was later strengthened by Rubinstein and Wolinsky (1994) by adding the requirement of common knowledge of rationality. Both the weaker and the stronger notions were further analyzed by Gilli (1999), while Esponda (2013) provided an explicit epistemic characterization.

would play R.[3] Note that in the game of Figure 2 there is no Nash equilibrium that yields the play aA. Thus a self-confirming equilibrium need not be a Nash equilibrium.

In the self-confirming equilibrium literature, a solution is still defined in terms of a *strategy profile*; in particular, a self-confirming equilibrium specifies choices at all unreached information sets. From a conceptual point of view, however, it is not clear what role choices at unreached information sets play beyond expressing the beliefs of the active players along the equilibrium path. For example, consider again the game of Figure 2 and a situation where Player 1 plays a believing that Player 3 would play L and Player 2 plays A believing that Player 3 would play R. Why is this not enough as a "solution"? Why the need to settle the counterfactual concerning what Player 3 would truly do if her information set were to be reached?[4] Furthermore, it is not clear how the counterfactual could be settled: both L and R can be justified as hypothetical rational choices for Player 3. The standard theory of counterfactuals, due to Robert Stalnaker and David Lewis (Stalnaker (1968), Stalnaker and Thomason (1970), Lewis (1973)), postulates a family of similarity relations on the set of possible worlds (one for each possible world) and the sentence "if ϕ were the case then ψ would be the case" is declared to be true at a possible world ω if ψ is true at the most similar world(s) to ω where ϕ is true. At a world where Players 1 and 2 play (a, A), we can take ϕ to be the sentence "Player 3's information set is reached" and ψ the sentence "Player 3 plays L". Then the sentence "if ϕ were the case then ψ would be the case" would be true at the actual world (where Player 3's information set is *not* reached, because Players 1 and 2 play (a, A)) if and only if the most similar world to the actual world is one where Player 3 plays L. But how are we to determine if the most similar world to the actual world is one where Player 3 plays L or one where Player 3 plays R?

In general, pinning down the counterfactual choices that would be made at unreached information sets is not a straightforward matter. Consider, for example, the game illustrated in Figure 3 due to Perea (Perea 2010, p. 169), where

[3] As Greenberg (2000) points out, Player 3 gains from the uncertainty in the minds of Players 1 and 2 and, if asked what she would do, she would refuse to answer, since her payoff is largest when Players 1 and 2 play aA.

[4] One answer could be: because we want to know which player has false beliefs about Player 3. But in what way would knowing this be useful? After all, presumably Players 1 and 2 are very confident of what they believe and the falsity of their beliefs can only be judged from an external point of view. Another answer could be: because we want to determine what would happen if either Player 1 made a mistake and played b instead of a or Player 2 made a mistake and played B instead of A. But why should a rationality-based solution concept be built on the possibility of players making mistakes?

Player 1 can either play b and end the game or play a in which case she and Player 2 play a "simultaneous" game. Using backward-induction reasoning,[5]

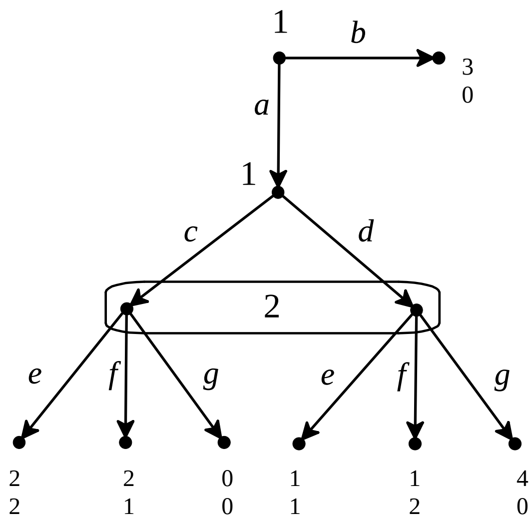

Figure 3: The conflict between backward induction and forward induction.

one first applies the procedure of iterative deletion of strictly dominated strategies to the subgame that starts at history a to obtain (c, e) and then one concludes that Player 1 will play b.[6] Thus the "rational solution" would be the strategy profile (b, e). On the other hand, applying forward-induction reasoning,[7] one first eliminates Player 1's strategy ac (since it is strictly dominated by b) and Player 2's strategy g; then one eliminates Player 1's strategy ad and Player 2's strategy e, with the conclusion that Player 1 will play b. Thus the "rational solution" would be the strategy profile (b, f). Note that the prediction in terms

[5] As captured, for example, by the notion of common belief in present and future rationality (Perea (2014)), or forward belief in rationality (Baltag et al. (2009), Bonanno (2014)). See also Penta (2009).

[6] In the subgame, for Player 2 g is strictly dominated (by both e and f); after deleting g, for Player 1 d becomes strictly dominated by c; after deleting d, for Player 2 f becomes strictly dominated by e.

[7] As captured, for example, by the notion of extensive-form rationalizability (Pearce (1984)), which is characterized by common strong belief in rationality (Battigalli and Siniscalchi (2002)).

of play is the same, namely that Player 1 would end the game by playing b, but the answer to the counterfactual question "what would Player 2 do at her unreached information set?" is different. Is it essential – or even important – that we provide an answer to this counterfactual question? Or can we rely on a *purely behavioral* theory of rational play, whose aim is merely to predict the *actual* behavior of rational players and not what would happen at unreached information sets?

In this paper we outline the elements of such a theory. The framework is rich enough to enable one to explicitly model the *actual* choices of the players and the beliefs that justify those choices. Thus we are aiming for an epistemic, or rather doxastic,[8] theory of rational behavior in extensive-form games. In the next section we elaborate on the distinction between a strategy-based approach and a behavior-based approach.

2 Strategy-based models *versus* behavioral models

The aim of the so-called *epistemic foundation program* in game theory is to characterize, for any game, the choices made by rational players who know the structure of the game and the preferences of their opponents and who recognize each other's rationality. The expression 'mutual recognition of rationality' has been interpreted in the literature as 'common belief in rationality'. Earlier notions – such as rationalizability (Bernheim (1984), Pearce (1984)) – captured only informally the concept of common belief in rationality. The first explicitly epistemic analysis is due to Aumann (1987), who provided an epistemic characterization of the notion of correlated equilibrium and Tan and Werlang (1988) who provided a doxastic characterization of rationalizability.[9]

There are two types of epistemic/doxastic models used in the game-theoretic literature: the so-called "state-space" models and the "type-space" models. We will adopt the former.[10]

In the standard state-space model of a given game, one takes as starting point a set of states (or possible worlds) and associates with every state a strategy for every player. If ω is a state and s_i is the strategy of player i at ω

[8] 'Epistemic' means 'knowledge-based', while 'doxastic' means 'belief-based'. The crucial distinction between knowledge and belief is that what is known must be true, while beliefs can be erroneous. From a conceptual point of view, it is essential to allow for the possibility of erroneous beliefs (see (Stalnaker 1996, p.153)).

[9] Other notable early contributions are Brandenburger and Dekel (1987), Stalnaker (1994; 1996). For an overview see Battigalli and Bonanno (1999), Perea (2012), Bonanno (2015a)

[10] Note that there is a straightforward way of translating one type of model into the other.

then the interpretation is that, at that state, player i plays s_i. If the game is simultaneous (so that there cannot be any unreached information sets), then there is no ambiguity in the expression "player i plays s_i", but if the game is an extensive-form game then the expression *is* ambiguous. Consider, for example, the game of Figure 1 and a state ω where both Players 1 and 2 play a, so that Player 3's information set is not reached. Suppose that the strategy of Player 3 associated with state ω is R. In what sense does Player 3 "play" R? Does it mean that, before the game starts, Player 3 has made a plan to play R if her information set were to be reached? Or does it mean (in a Stalnaker-Lewis interpretation of the counterfactual) that in the state most similar to ω where her information set is actually reached, Player 3 plays R?[11] Or is Player 3's strategy R to be interpreted not as a statement about what Player 3 would actually do but as an expression of what each of her opponents thinks that Player 3 would do?[12]

While most of the literature on the epistemic foundations of game theory makes use of strategy-based models, a few papers[13] follow a *behavioral* approach by associating with each state a *play* (or outcome) of the game. The challenge in this class of models is to capture the reasoning of a player who takes a particular action while considering what would happen if she took a different action. The most common approach is to postulate, for each player, a set of conditional beliefs, where the conditioning events are represented by possible histories in the game (including off-path histories).[14] Here we will follow the simpler approach put forward in Bonanno (2013a; 2014; 2018), which models the "deliberation-stage" beliefs of a player. Previously, the literature captured the "after-choice" beliefs; in particular, it was based on the assumption that if at a state a player takes action a then the player knows that she takes action a. The deliberation-stage approach, on the other hand, models the beliefs of the player at the time when she is considering all the actions available to her and treats each of those actions as "open possibilities". Thus her beliefs take the

[11] Halpern (2001) adopts this interpretation after pointing out that in this type of models "one possible culprit for the confusion in the literature regarding what is required to force the backward induction solution in games of perfect information is the notion of a strategy".

[12] The conceptual issues that arise in strategy-based models of extensive-form games are discussed at length in Bonanno (2015b).

[13] The seminal contribution is Samet (1996), followed by Baltag et al. (2009), Battigalli et al. (2013), Bonanno (2013a; 2014; 2018). With the exception of Bonanno (2014), this literature has focused on games with perfect information.

[14] Samet (1996) uses extended information structures to model hypothetical knowledge (see Halpern (1999) for a critical discussion of extended information structures), Baltag et al. (2009) use plausibility relations and Battigalli et al. (2013) use conditional probability systems.

following form: "if I take action a then the outcome will be x and if I take action b then the outcome will be y". In this approach the conditional "if p then q" is interpreted as a *material conditional*, that is, as equivalent to "either not p or q".[15] Hence this analysis does not rely in any way on (objective or subjective) counterfactuals; furthermore, only the beliefs of the active players at the time of choice are modeled: *no initial beliefs nor belief revision policies are postulated*. The approach is described in the next section, after first recalling the history-based definition of extensive-form game[16] and the notion of qualitative belief.

3 Behavioral models of extensive-form games

3.1 The history-based definition of extensive-form game

For simplicity we will restrict attention to games with ordinal payoffs and without chance moves. We will *not*, however, make the common assumption of "no relevant ties" or genericity of payoffs; furthermore we allow for imperfect information.

If A is a set, we denote by A^* the set of finite sequences in A. If $h = \langle a_1, ..., a_k \rangle \in A^*$ and $1 \leq i \leq k$, the sequence $h' = \langle a_1, ..., a_i \rangle$ is called a *prefix* of h and we denote this by $h' \precsim h$; furthermore, if $h' \precsim h$ and $h' \neq h$ then we write $h' \prec h$ and say that h' is a *proper prefix* of h. If $h = \langle a_1, ..., a_k \rangle \in A^*$ and $a \in A$, we denote the sequence $\langle a_1, ..., a_k, a \rangle \in A^*$ by ha.

A *finite extensive form without chance moves* is given by the following elements:

1. A finite set N of players.

2. A finite set A of actions.

3. A finite set of histories $H \subseteq A^*$ which is closed under prefixes (that is, if $h \in H$ and $h' \in A^*$ is a prefix of h, then $h' \in H$). The null history $\langle\rangle$, denoted by $\emptyset$, is an element of H and is a prefix of every history. A history $h \in H$ such that, for every $a \in A$, $ha \notin H$, is called a *terminal history* or *play*. The set of terminal histories is denoted by Z. $D = H \backslash Z$ denotes the set of non-terminal or *decision* histories.

[15] In Bonanno (2021) it is argued that, contrary to a common view, the material conditional is indeed sufficient to model deliberation.

[16] See, for example, Osborne and Rubinstein (1994).

4. A function $\iota : D \to N$ that assigns to each decision history h the player who moves, or *is active*, at h.[17] For every decision history $h \in D$, we denote by $A(h)$ the set of actions available at h (to player $\iota(h)$), that is, $a \in A(h)$ if and only if $a \in A$ and $ha \in H$.

5. For every player $i \in N$, an equivalence relation $\approx_i$ on D_i, where $D_i = \{h \in D : i = \iota(h)\}$ is the set of histories at which player i is active. The interpretation of $h \approx_i h'$ is that, when choosing an action at history $h \in D_i$, player i does not know whether she is moving at h or at h'. The equivalence class of $h \in D$ is denoted by $[h]$ and is called an *information set* of player $\iota(h)$; thus $[h] = \{h' \in D_{\iota(h)} : h \approx_{\iota(h)} h'\}$. The following restriction applies to information sets: if $h \approx_i h'$ then $A(h') = A(h)$, that is, the set of actions available to a player is the same at any two histories that belong to the same information set of that player. We also assume the property of *perfect recall*, according to which a player always remembers her own past moves: if $h_1, h_2 \in D_i$, $a \in A(h_1)$ and $h_1 a$ is a prefix of h_2 then, for every h' such that $h' \approx_i h_2$, there exists an $h \approx_i h_1$ such that ha is a prefix of h'.
   If every information set is a singleton (that is, if $h \approx_{\iota(h)} h'$ implies that $h = h'$, for every decision history h) then the game is said to have *perfect information*, otherwise it is said to have *imperfect information*.

From now on, non-null histories will be denoted succinctly by listing the corresponding actions, without brackets, without commas and omitting the empty history: thus instead of writing $\langle \emptyset, a_1, a_2, a_3, a_4 \rangle$ we will simply write $a_1 a_2 a_3 a_4$.

Given an extensive form, one obtains an *extensive game with ordinal payoffs* by adding, for every player $i \in N$, a complete and transitive preference relation R_i over the set Z of terminal histories (the interpretation of zR_iz' is that player i considers terminal history z to be at least as good as terminal history z'). It is customary to replace the preference relation R_i with a *utility* (or *payoff*) *function* $u_i : Z \to \mathbb{R}$ (where $\mathbb{R}$ denotes the set of real numbers) satisfying the property that $u_i(z) \geq u_i(z')$ if and only if $z\, R_i z'$.

We only consider ordinal payoffs and qualitative beliefs in order to highlight the important features of our approach in as simple a framework as possible. The analysis can be extended to the case where the players' preferences are represented by von Neumann-Morgenstern utility functions and beliefs are

[17] For notational simplicity we do not allow more than one player to be active at a decision history. A simultaneous move by, say, Players 1 and 2 is thus represented in the traditional way by having Player 1 move first followed by Player 2, who is not informed of Player 1's move.

probabilistic.[18]

3.2 Qualitative beliefs

Before introducing the definition of a model of a game, we recall the following facts about belief relations and operators.

Let Ω be a set, whose elements are called *states* (or possible worlds). We represent the beliefs of an agent by a binary relation $\mathcal{B} \subseteq \Omega \times \Omega$. The interpretation of $(\omega, \omega') \in \mathcal{B}$, also denoted by $\omega \mathcal{B} \omega'$, is that at state ω the agent considers state ω' possible; we also say that ω' is *reachable from* ω by $\mathcal{B}$. For every $\omega \in \Omega$ we denote by $\mathcal{B}(\omega)$ the set of states that are reachable from ω, that is, $\mathcal{B}(\omega) = \{\omega' \in \Omega : \omega \mathcal{B} \omega'\}$.

$\mathcal{B}$ is *transitive* if $\omega' \in \mathcal{B}(\omega)$ implies $\mathcal{B}(\omega') \subseteq \mathcal{B}(\omega)$ and it is *euclidean* if $\omega' \in \mathcal{B}(\omega)$ implies $\mathcal{B}(\omega) \subseteq \mathcal{B}(\omega')$. We will assume throughout that the belief relations are transitive and euclidean (see below for an interpretation of this assumption in terms of introspection).

Graphically, we represent a transitive and euclidean belief relation as shown in Figure 4, by adopting the following convention: for any two states $\omega, \omega' \in \Omega$, $\omega' \in \mathcal{B}(\omega)$ if and only if either ω and ω' are enclosed in the same rounded rectangle or there is an arrow from ω to the rounded rectangle containing ω'.[19]

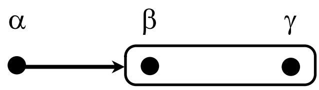

Figure 4: The relation $\mathcal{B} = \{(\alpha,\beta),(\alpha,\gamma),(\beta,\beta),(\beta,\gamma),(\gamma,\beta),(\gamma,\gamma)\}$.

[18] The traditional approach postulates that every player has a preference relation over the set of *lotteries* over terminal histories that satisfies the axioms of expected utility. This is not an innocuous assumption, since the game under consideration is implicitly taken to be common knowledge among the players. Thus not only is it commonly known who the players are, what choices they have available and what the possible outcomes are, but also how each player ranks those outcomes. While it is certainly reasonable to postulate that a player knows her own preferences, it is much more demanding to assume that she knows the preferences of her opponents. If those preferences are expressed as ordinal rankings, this assumption is less troublesome than in the case where preferences also incorporate attitudes to risk (that is, the utility functions that represent those preferences are von Neumann-Morgenstern utility functions).

[19] In other words, for any two possible worlds ω and ω' that are enclosed in a rounded rectangle, $\{(\omega,\omega),(\omega,\omega'),(\omega',\omega),(\omega',\omega')\} \subseteq \mathcal{B}$ and, if there is an arrow from a state ω to a rounded rectangle, then, for every ω' in that rectangle, $(\omega,\omega') \in \mathcal{B}$.

The object of beliefs are propositions or events (i.e. sets of states; events are denoted by bold-type capital letters). We say that *at state ω the agent believes event* $\mathbf{E} \subseteq \Omega$ if and only if $\mathcal{B}(\omega) \subseteq \mathbf{E}$. For example, in the case illustrated in Figure 4, at state α the agent believes event $\{\beta, \gamma\}$. If $\mathbf{E}$ is an event then the sentence "the agent believes $\mathbf{E}$" can itself be interpreted as an event, denoted by $\mathbb{B}\mathbf{E} \stackrel{\text{def}}{=} \{\omega \in \Omega : \mathcal{B}(\omega) \subseteq \mathbf{E}\}$. For example, in the case illustrated in Figure 4, $\mathbb{B}\{\beta, \gamma\} = \{\alpha, \beta, \gamma\}$. Thus one can define a *belief operator* $\mathbb{B} : 2^{\Omega} \to 2^{\Omega}$ that associates with every event $\mathbf{E}$ the (possibly empty) event $\mathbb{B}\mathbf{E}$ that the individual believes $\mathbf{E}$. It is well known that transitivity of the relation $\mathcal{B}$ corresponds to positive introspection of beliefs (if the individual believes $\mathbf{E}$ then she believes that she believes $\mathbf{E}$: $\mathbb{B}\mathbf{E} \subseteq \mathbb{B}\mathbb{B}\mathbf{E}$) and euclideanness corresponds to negative introspection (if the individual does not believe $\mathbf{E}$ then she believes that she does not believe $\mathbf{E}$: $\neg\mathbb{B}\mathbf{E} \subseteq \mathbb{B}\neg\mathbb{B}\mathbf{E}$; for every event $\mathbf{F}$, $\neg\mathbf{F}$ denotes the complement of $\mathbf{F}$ in Ω).[20]

We say that, *at state ω, event* $\mathbf{E}$ *is true* if $\omega \in \mathbf{E}$. In the case illustrated in Figure 4, at state α the agent erroneously believes event $\{\beta, \gamma\}$: $\alpha \in \mathbb{B}\{\beta, \gamma\}$ but the event $\{\beta, \gamma\}$ is not true at α; this is due to the fact that $\alpha \notin \mathcal{B}(\alpha)$. We say that *at state ω the agent has correct beliefs* if $\omega \in \mathcal{B}(\omega)$. Thus we can define the event that the agent has correct beliefs, denoted by $\mathbf{T}^*$ (where 'T' stands for 'true'), as follows:

$$\mathbf{T}^* = \{\omega \in \Omega : \omega \in \mathcal{B}(\omega)\}. \tag{1}$$

For example, in the case illustrated in Figure 4, $\mathbf{T}^* = \{\beta, \gamma\}$.

Remark 1. *Note that it is a consequence of euclideanness of the relation $\mathcal{B}$ that* the agent always believes that her beliefs are correct*: if $\omega' \in \mathcal{B}(\omega)$ then $\omega' \in \mathcal{B}(\omega')$ or, equivalently, $\mathbb{B}\mathbf{T}^* = \Omega$.*

3.3 Behavioral models of games

To define a model of a game, we begin with a set of states Ω and interpret each state in terms of a particular play of the game by means of a function $\zeta : \Omega \to Z$ that associates, with every state ω, a terminal history $\zeta(\omega) \in Z$. Next we add, for every decision history $h \in D$, a binary relation $\mathcal{B}_h$ on Ω representing the beliefs of $\iota(h)$, the active player at h;[21] however, we do so only at histories that

[20] For more details see Battigalli and Bonanno (1999).

[21] It would be more precise to write $\mathcal{B}_{\iota(h)}$ instead of $\mathcal{B}_h$, but we have chosen the lighter notation since there is no ambiguity, because we have assumed that at every decision history there is a unique player who is active there.

are actually reached at a given state, in the sense that $\mathcal{B}_h(\omega) \neq \varnothing$ if and only if $h \prec \zeta(\omega)$.

Definition 3.1. Given an extensive-form game, a *model of it* is a tuple $\langle \Omega, \zeta, \{\mathcal{B}_h\}_{h \in D} \rangle$ where

- Ω is a set of states.
- $\zeta : \Omega \to Z$.
- For every $h \in D$, $\mathcal{B}_h \subseteq \Omega \times \Omega$ is a belief relation that satisfies the following properties:
  1. $\mathcal{B}_h(\omega) \neq \varnothing$ if and only if $h \prec \zeta(\omega)$ [beliefs are specified only at reached decision histories and are consistent].[22]
  2. If $\omega' \in \mathcal{B}_h(\omega)$ then $h' \prec \zeta(\omega')$ for some h' such that $h' \approx_{\iota(h)} h$ [the active player at history h correctly believes that her information set that contains h has been reached].[23]
  3. If $\omega' \in \mathcal{B}_h(\omega)$ then (1) $\mathcal{B}_h(\omega') = \mathcal{B}_h(\omega)$ and (2) if $h' \prec \zeta(\omega')$ with $h' \approx_{\iota(h)} h$ then $\mathcal{B}_{h'}(\omega') = \mathcal{B}_h(\omega)$ [by (1), beliefs satisfy positive and negative introspection and, by (2), beliefs are the same at any two histories in the same information set; thus one can unambiguously refer to a player's beliefs at an information set].
  4. If $\omega' \in \mathcal{B}_h(\omega)$ and $h' \prec \zeta(\omega')$ with $h' \approx_{\iota(h)} h$, then, for every action $a \in A(h)$, there is an $\omega'' \in \mathcal{B}_h(\omega)$ such that $h'a \precsim \zeta(\omega'')$.

The last condition states that if, at state ω and history h reached at ω ($h \prec \zeta(\omega)$), player $\iota(h)$ considers it possible that the play of the game has reached history h', which belongs to her information set that contains h, then, for *every* action a available at that information set, there is a state ω'' that she considers possible at h and ω ($\omega'' \in \mathcal{B}_h(\omega)$) where she takes action a at history h' ($h'a \precsim \zeta(\omega'')$). This means that, for every available action, the active player at h has a belief about what will happen if she chooses that action.[24] A further "natural" restriction on beliefs will be discussed later (Section 4).

[22]Consistency means that there is no event **E** such that both **E** and $\neg$**E** are believed. It is well known that, at state ω, beliefs are consistent if and only if $\mathcal{B}(\omega) \neq \varnothing$.

[23]Recall that $h' \approx_{\iota(h)} h$ (also written as $h' \in [h]$) if and only if h and h' belong to the same information set of player $\iota(h)$ (thus $\iota(h) = \iota(h')$).

[24]As noted above, this way of modeling beliefs is a departure from the standard approach in the literature, where it is assumed that if, at a state, a player takes a particular action then she knows that she takes that action. The standard approach thus requires the use of either objective or

Figure 5 reproduces the game of Figure 1 and shows a model of it. For every history, under each state that is considered possible (by the corresponding player) we have recorded the action actually taken by the player at that state and the player's payoff (at the terminal history associated with that state).

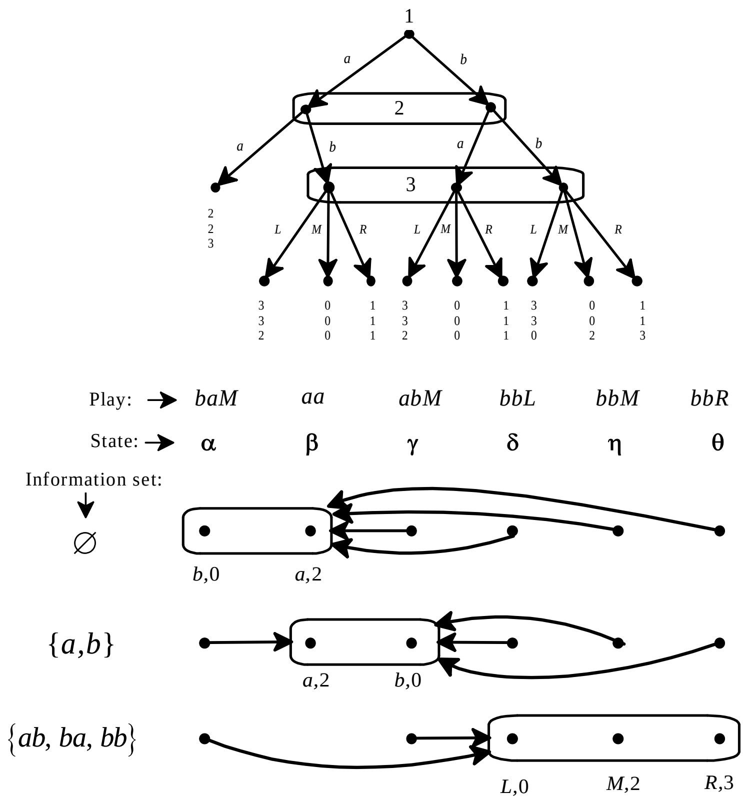

Figure 5: The game of Figure 1 and a model of it.

Suppose, for example, that the actual state is γ. State γ encodes the following.

subjective counterfactuals in order to represent a player's beliefs about the consequences of taking alternative actions. In our approach a player's beliefs refer to the *deliberation* or *pre-choice* stage, where the player considers the consequences of taking *any* available action, without pre-judging her subsequent decision.

1. As a matter of fact, Player 1 plays a, Player 2 plays b and Player 3 plays M.

2. Player 1 (who is active at history $\emptyset$) believes that if she plays a then Player 2 will also play a (this belief is erroneous since at state γ Player 2 actually plays b, after Player 1 plays a) and thus her utility will be 2, and she believes that if she plays b then Player 2 will play a and Player 3 will play M and thus her utility will be 0.

3. Player 2 (who is active at information set $\{a, b\}$) correctly believes that Player 1 played a and, furthermore, correctly believes that if he plays b then Player 3 will play M and his utility will be 0, and believes that if he plays a his utility will be 2.

4. Player 3 (who is active at information set $\{ab, ba, bb\}$) erroneously believes that both Player 1 and Player 2 played b; thus, she believes that if she plays L her utility will be 0, if she plays M her utility will be 2 and if she plays R her utility will be 3.

On the other hand, if the actual state is β, then the actual play is aa and the beliefs of Players 1 and 2 are as detailed above (points 2 and 3, respectively), while *no beliefs are specified for Player 3*, because Player 3 does not get to play (her information set is not reached).

3.4 Rationality

Consider again the model of Figure 5 and state γ. There Player 1 believes that if she takes action a then her utility will be 2 and if she takes action b then her utility will be 0. Thus, if she is rational, she must take action a. Indeed, at state γ she does take action a and thus she is rational (although she will later discover that her belief was erroneous and her utility turns out to be 0). Since Player 1 has the same beliefs at every state, we declare Player 1 to be rational at precisely those states where she takes action a, namely β and γ. Similar reasoning leads us to conclude that Player 2 is rational at those states where she takes action a, namely states α and β. Similarly, Player 3 is rational at state θ only. If we denote by $\mathbf{R}$ the event that all the *active* players are rational, then in the model of Figure 5 we have that $\mathbf{R} = \{\beta\}$ (note that at state β Player 3 is *not* active).

We need to define the notion of rationality more precisely. Various definitions of rationality have been suggested in the context of extensive-form games, most notably *material rationality* and *substantive rationality* (Aumann

(1995; 1998)). The former notion is weaker in that a player can be found to be irrational only at decision histories of hers that are actually reached. The latter notion, on the other hand, is more stringent since a player can be judged to be irrational at a decision history h of hers even if it is not reached. Given that we have adopted a purely behavioral approach, the natural notion for us is material rationality. We will adopt a very weak version of it, according to which at a state ω and reached history h (that is, $h \prec \zeta(\omega)$), the active player at h is rational if the following is the case: if a is the action that the player takes at h at state ω (that is, $ha \precsim \zeta(\omega)$) then there is no other action at h that, according to her beliefs, *guarantees* a higher utility.

Definition 3.2. Let ω be a state, h a decision history that is reached at ω ($h \prec \zeta(\omega)$) and $a, b \in A(h)$ two actions available at h.

(A) We say that, at ω and h, the active player $\iota(h)$ *believes that b is better than a* if, for all $\omega_1, \omega_2 \in \mathcal{B}_h(\omega)$ and for every h' such that $h' \approx_{\iota(h)} h$ (that is, history h' belongs to the same information set as h), if $h'a \precsim \zeta(\omega_1)$ (that is, a is the action taken at history h' at state ω_1) and $h'b \precsim \zeta(\omega_2)$ (that is, b is the action taken at h', at state ω_2) then $u_{\iota(h)}(\zeta(\omega_1)) < u_{\iota(h)}(\zeta(\omega_2))$. In other words, the active player at h believes that b is better than a if, *restricting attention to the states that she considers possible*, the *maximum* utility that she obtains if she plays a is less than the *minimum* utility that she obtains if she plays b.

(B) We say that player $\iota(h)$ is *rational at history h at state ω* if and only if the following is true: if $ha \precsim \zeta(\omega)$ (that is, $a \in A(h)$ is the action played at h at state ω) then, for every $b \in A(h)$, it is not the case that, at state ω and history h, player $\iota(h)$ believes that b is better than a.

Finally, we define the event that all the *active* players are rational, denoted by $\mathbf{R}$ as follows:

$$\begin{array}{c}\omega \in \mathbf{R} \text{ if and only if, for every } h \prec \zeta(\omega), \\ \text{player } \iota(h) \text{ is rational at } h \text{ (at state } \omega). \end{array} \tag{2}$$

For example, as noted above, in the model of Figure 5 we have that $\mathbf{R} = \{\beta\}$.

3.5 Correct beliefs

The notion of locally correct belief was introduced in Section 3.2 and was identified with *local* reflexivity (that is, reflexivity at a state, rather than global reflexivity). Since, at any state, only the beliefs of the active players are specified, we define the event that players have correct beliefs by restricting attention to those players who actually move. Thus the event that the *active players have correct beliefs*, denoted by **T**, is defined as follows:

$$\omega \in \mathbf{T} \text{ if and only if } \omega \in \mathcal{B}_h(\omega) \text{ for every } h \text{ such that } h \prec \zeta(\omega). \qquad (3)$$

For example, in the model of Figure 5, $\mathbf{T} = \{\beta\}$.

What does the expression "correct beliefs" mean? Consider state β in the model of Figure 5 where the active players (Players 1 and 2) have correct beliefs in the sense of (3) ($\beta \in \mathcal{B}_\emptyset(\beta)$ and $\beta \in \mathcal{B}_a(\beta)$). Consider Player 1. There are two components to Player 1's beliefs: (*i*) she believes that if she plays a then Player 2 will also play a, (*ii*) she believes that if she plays b then Players 2 and 3 will play a and M, respectively. The first belief is correct *at state* β, where Player 1 plays a and Player 2 indeed follows with a. As for the second belief, whether it is correct or not depends on how we interpret it. If we interpret it as the material conditional "if b then (a, M)" (which is equivalent to "either not b or (a, M)") then it is indeed true at state β, but trivially so, because the antecedent is false there (Player 1 does *not* play b). If we interpret it as a counterfactual conditional "if Player 1 *were* to play b then Players 2 and 3 *would* play a and M, respectively" then in order to decide whether the conditional is true or not one would need to enrich the model by adding a "similarity" or "closeness" relation on the set of states (in the spirit of Stalnaker (1968), Lewis (1973)); one would then check if at the closest state(s) to β at which Player 1 plays b it is indeed the case that Players 2 and 3 play a and M, respectively. Note that there is no *a priori* reason to think that the closest state to β is α. This is because, as pointed out by Stalnaker (Stalnaker 1998, p.48), there is no necessary connection between counterfactuals, which capture causal relations, and beliefs: for example, I can believe that, if I drop the vase that I am holding in my hands, it will break (because I believe it is made of glass) but in fact, if I were to drop it, it would not break (because it is actually made of plastic).

Our models do not have the resources to answer the question: "*at state* β, is it true – as Player 1 believes – that if Player 1 were to play b then Players 2 and 3 would play a and M, respectively?" One could, of course, enrich the models, but is there a compelling reason to do so? In other words, is it important to be able to answer such questions? If we are merely interested in determining what rational players *do*, then what matters is what actions they actually take and what they believe when they act, whether or not those beliefs are correct in a stronger sense than is captured by the material conditional.

Is the material conditional interpretation of "if I play a then the outcome will be x" sufficient, though? Since the crucial assumption in the proposed framework is that the agent considers all of her available actions as possible (that is, for every available action there is a doxastically accessible state where she takes that action), material conditionals are indeed sufficient: the material conditional "if I take action a the outcome will be x" zooms in – through the lens of the agent's beliefs – on those states where action a is indeed taken and verifies that at those states the outcome is indeed x, while the worlds where action a is not taken are not relevant for the truth of the conditional.[25]

3.6 Self-confirming play

We have defined two events: the event **R** that all the *active* players are rational and the event **T** that all the *active* players have locally correct beliefs. In the model of Figure 5 we have that $\mathbf{R} \cap \mathbf{T} = \{\beta\}$ and it so happens that $\zeta(\beta) = aa$ is a Nash equilibrium play, that is, there is a pure-strategy Nash equilibrium whose associated play is aa (in fact, there are two: (a, a, M) and (a, a, R)). However, as shown below, this is not always the case.

At a play associated with a state $\omega \in \mathbf{R} \cap \mathbf{T}$, each active player's chosen action is rationally justified by her beliefs at the time of choice (since $\omega \in \mathbf{R}$) and the beliefs concerning what would happen after that action turn out to be correct (since $\omega \in \mathbf{T}$), so that no player is faced with evidence that her beliefs were wrong. Does that mean that, once the final outcome $\zeta(\omega)$ is revealed, no player regrets her actual choice? The answer is negative, because it is possible that a player, while not having any false beliefs, might not anticipate with precision the actions of the players who move after her. In the model shown in Figure 6 we have that $\mathbf{R} \cap \mathbf{T} = \{\alpha, \beta, \gamma\}$, that is, at every state the active players are rational and have correct beliefs. Consider state β, where the play is ad. At state β Player 1 is rational because she believes that if she plays b her utility will be 1 and if she plays a her utility might be 0 but might also be 2 (she is uncertain about what Player 2 will do). Thus she does not believe that action b is better than a and hence it is rational for her to play a (Definition 3.2). Player 2 is rational because she is indifferent between her two actions. However, *ex post*, when Player 1 learns that the actual outcome is ad, she regrets not taking action b instead of a. This example shows that even though $\beta \in \mathbf{R} \cap \mathbf{T}$, $\zeta(\beta) = ad$ is not a Nash equilibrium play.

Next we introduce another event which, in conjunction with **T**, guarantees that the active players' beliefs about the opponents' actual moves are exactly

[25] For a more in-depth discussion see Bonanno (2021).

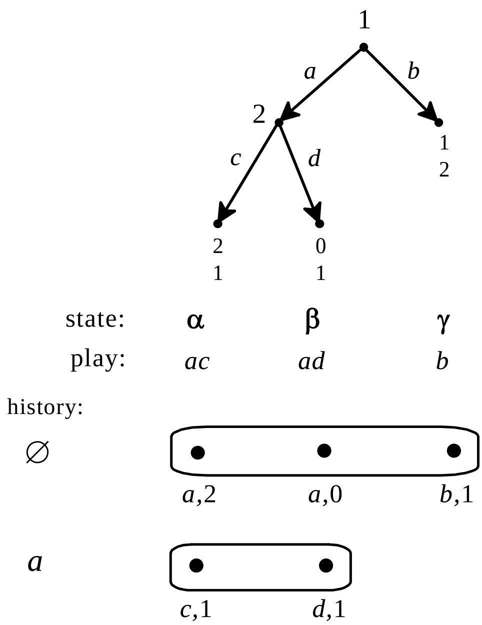

Figure 6: A perfect-information game and a model of it.

correct.[26] Event **C** ('C' for 'certainty') defined below rules out uncertainty about the opponents' *past* choices (Point 1) as well as uncertainty about the opponents' *future* choices (Point 2). Note that Point 1 is automatically satisfied in games with perfect information and thus imposes restrictions on beliefs only in imperfect-information games.

Definition 3.3. A state ω belongs to event **C** if and only if, for every reached history at ω (that is, for every $h \prec \zeta(\omega)$), and $\forall \omega', \omega'' \in \mathcal{B}_h(\omega)$, $\forall h', h'' \in [h]$ (recall that $[h]$ is the information set that contains h),

1. if $h' \prec \zeta(\omega')$ and $h'' \prec \zeta(\omega'')$ then $h' = h''$,

2. $\forall a \in A(h)$, if $h'a \prec \zeta(\omega')$ and $h''a \prec \zeta(\omega'')$ then $\zeta(\omega') = \zeta(\omega'')$.

[26] Note that a requirement built in the definition of a self-fulfilling equilibrium (Fudenberg and Levine 1993, p.523) is that "each player's beliefs about the opponents' play are exactly correct".

Note that – concerning Point 1 – a player may be erroneous in her certainty about the opponents' past choices (that is, it may be that $\omega \in \mathbf{C}$, the actual reached history is $h \prec \zeta(\omega)$ and yet player $\iota(h)$ is certain that she is moving at history $h' \in [h]$ with $h' \neq h$)[27] and – concerning Point 2 – a player may also be erroneous in her certainty about what will happen after her choice.[28]
In the model of Figure 5 $\mathbf{C} = \Omega$, while in the model of Figure 6 $\mathbf{C} = \varnothing$, because at history $\varnothing$ Player 1 is uncertain about what will happen if she takes action a.

If state ω belongs to the intersection of events $\mathbf{C}$ and $\mathbf{T}$ then, at state ω, each active player's beliefs about the opponents' *actual* play are exactly correct. Note, however, that – as noted in Section 3.5 – there is no way of telling whether or not a player is also correct about what would happen after her counterfactual choices, because the models that we are considering are not rich enough to address the issue of counterfactuals.

Definition 3.4. Let G be a game and z a play (or terminal history) in G. We say that z is a *self-confirming play* if there exists a model of G and a state ω in that model such that (1) $\omega \in \mathbf{R} \cap \mathbf{T} \cap \mathbf{C}$ and (2) $z = \zeta(\omega)$.

Definition 3.5. Given a game G and a play z in G, call z a *Nash play* if there is a pure-strategy Nash equilibrium whose induced play is z.

It turns out that, in perfect-information games in which no player moves more than once along any play, the two notions of self-confirming play and Nash play are equivalent.[29]

Proposition 1 (Bonanno (2018)). *Consider a perfect-information game G where no player moves more than once along any play. Then,*

(A) if terminal history z is a Nash play of G then there is a model of G and a state ω in that model such that (1) $\omega \in \mathbf{R} \cap \mathbf{T} \cap \mathbf{C}$ and (2) $\zeta(\omega) = z$,

(B) for any model of G and for every state ω in that model, if $\omega \in \mathbf{R} \cap \mathbf{T} \cap \mathbf{C}$ then $\zeta(\omega)$ is a Nash play.

[27] For example, in the model of Figure 5, at state γ (which belongs to event $\mathbf{C}$) and reached history ab, Player 3 is certain that she is moving at history bb while, as a matter of fact, she is moving at history ab.

[28] For example, in the model of Figure 5, at state γ and history $\varnothing$, Player 1 is certain that if she takes action a then Player 2 will also play a, but she is wrong about this, because, as a matter of fact, at state γ Player 2 follows with b rather than a.

[29] Note that every finite perfect-information game has at least one pure-strategy Nash equilibrium. Note also that, unlike most of the literature, we do not restrict attention to games with no relevant ties or generic payoffs.

For games with imperfect information, while it is still true that a Nash play is a self-confirming play, there may be self-confirming plays that are not Nash plays. That is, while Part (A) of Proposition 1 is true also for imperfect-information games, Part (B) is not. The reason for this is that two players might have different beliefs about the potential choice of a third player. Figure 7 reproduces the game of Figure 2 together with a model of it. In the model of

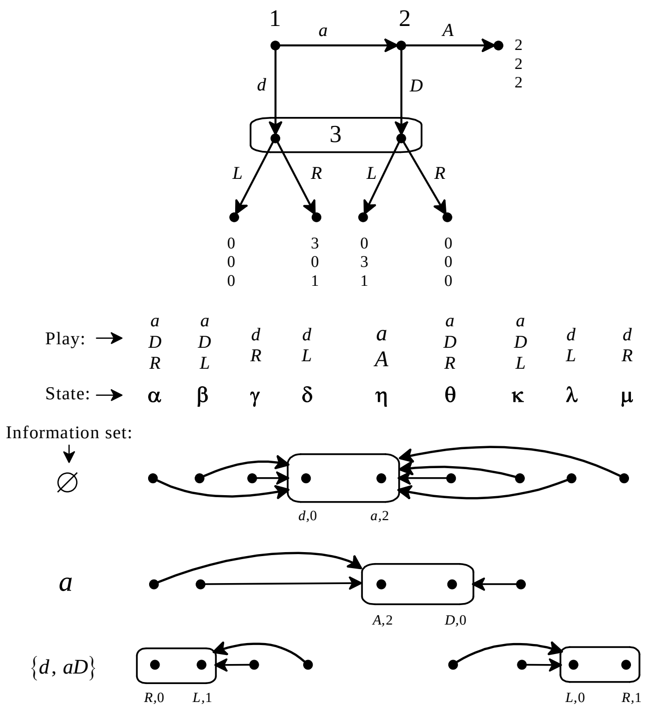

Figure 7: The game of Figure 2 and a model of it.

Figure 7, $\mathbf{R} = \mathbf{T} = \{\eta\}$ and $\mathbf{C} = \Omega$, so that $\mathbf{R} \cap \mathbf{T} \cap \mathbf{C} = \{\eta\}$. Thus at state η the active players (Players 1 and 2) are rational, have correct beliefs and have no uncertainty and yet $\zeta(\eta) = aA$ which is not a Nash play. Players 1 and 2 have different beliefs about what Player 3 would do at her information set: at state η Player 1 believes that if she plays d then Player 3 will play L, while Player 2 believes that if she plays D then Player 3 will play R. In the next section we introduce a new event, denoted by **A** ('A' stands for 'agreement'), that rules out such disagreement.

4 Nash play

In this section we provide a doxastic characterization of Nash play in general games (with possibly imperfect information). First we need to add one more condition to the definition of a model of a game that is relevant only if the game has imperfect information.

The definition of model given in Section 3.3 (Definition 3.1) allows for "causally flawed" beliefs. For example, consider a game where Player 1 moves first, choosing between a and b, and Player 2 moves second and chooses between c and d *without being informed of Player 1's choice* (that is, histories a and b belong to the same information set of Player 2). Definition 3.1 allows Player 1 to have the following beliefs: "if I play a, then Player 2 will play c, while if I play b then Player 2 will play d". Such beliefs can be considered "irrational" on the grounds that there cannot be a causal link between Player 1's move and Player 2's choice, since Player 2 does not get to observe Player 1's move.[30] Thus a "causally correct" belief for Player 1 would require that the predicted choice(s) of Player 2 be the same, no matter what action Player 1 herself chooses.

Definition 4.1. A *causally restricted model* of a game is a model (Definition 3.1) that satisfies the following additional restriction (a verbal interpretation follows).[31]

5. Let ω be a state, h a decision history reached at ω ($h \prec \zeta(\omega)$) and a and b two actions available at h ($a, b \in A(h)$). Let h_1 and h_2 be two decision histories that belong to the same information set of player $j = \iota(h_1)$ ($h_1 \approx_j h_2$) and c_1, c_2 be two actions available at h_1 ($c_1, c_2 \in A(h_1) = A(h_2)$). Then the following holds (recall that $[h]$ denotes the information set that contains decision history h, that is, $h' \in [h]$ if and only if $h' \approx_{\iota(h)} h$):

$$\begin{array}{l}\text{if } h', h'' \in [h],\ \omega_1, \omega_2 \in \mathcal{B}_h(\omega),\ h'a \prec h_1c_1 \precsim \zeta(\omega_1) \text{ and} \\ h''b \prec h_2c_2 \precsim \zeta(\omega_2), \text{ then either } c_1 = c_2 \text{ or there exist} \\ \omega_1', \omega_2' \in \mathcal{B}_h(\omega) \text{ such that } h'a \prec h_1c_2 \precsim \zeta(\omega_1') \text{ and} \\ h''b \prec h_2c_1 \precsim \zeta(\omega_2').\end{array} \tag{4}$$

[30]Several authors, however, have argued that such beliefs are not necessarily irrational: see, for example, Nozick (1969), Gauthier (1986), Bicchieri and Green (1999), Spohn (2003; 2007; 2010).

[31]Note that, for games with perfect information, there is no difference between a model and a restricted model, since (4) is vacuously satisfied.

In words: if, at state ω and reached history h, player $i = \iota(h)$ considers it possible that, if she takes action a, history h_1 is reached and player $j = \iota(h_1)$ takes action c_1 at h_1 and player i also considers it possible that, if she takes action b, then history h_2 is reached, which belongs to the same information set as h_1, and player j takes action c_2 at h_2, then either $c_1 = c_2$ or at state ω and history h player i must also consider it possible that (1) after taking action a, h_1 is reached and player j takes action c_2 at h_1 and (2) after taking action b, h_2 is reached and player j takes action c_1 at h_2.

Figure 8 shows a game and four partial models of it, giving only the beliefs of Player 1 (at history $\varnothing$): two of them violate Condition 5 of Definition 4.1 (the ones on the left that are labeled "not allowed"), while the other two satisfy it. Note that the models shown in Figures 5-7 are causally restricted models.

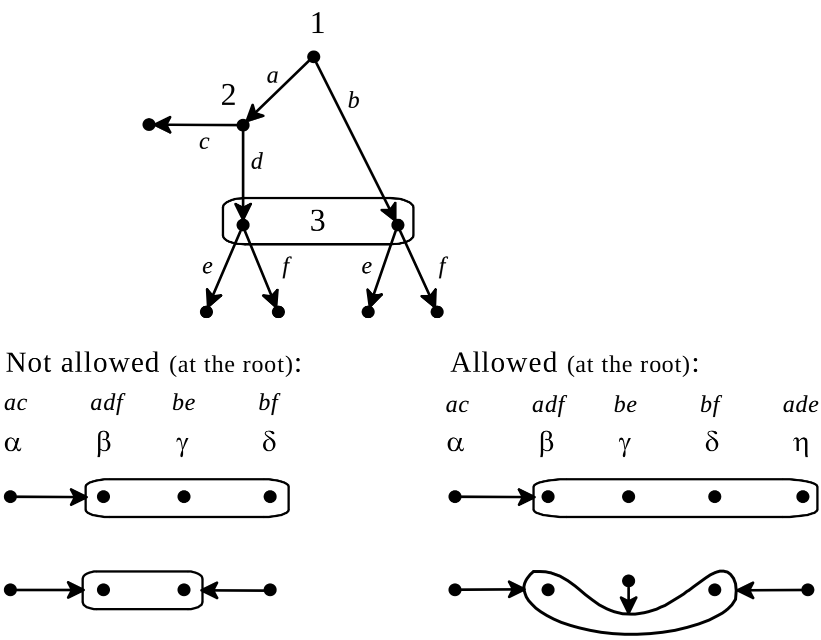

Figure 8: A game and four partial models of it (showing only the beliefs of Player 1 at history $\varnothing$), two of which violate Condition 5 of Definition 4.1 and the remaining two do not.

Now we turn to the notion of agreement, which is intended to rule out situations like the one shown in Figure 7 where Players 1 and 2 disagree about what action Player 3 would take at her information set $\{d, aD\}$.

Definition 4.2. We say that at state ω active players i and j *consider future information set* $[h]$ of player $k = \iota(h)$ if there exist

1. two decision histories h_1 and h_2 that are reached at ω (that is, $h_1 \prec h_2 \prec \zeta(\omega)$) and belong to i and j, respectively (that is, $i = \iota(h_1)$ and $j = \iota(h_2)$),

2. states $\omega_1 \in \mathcal{B}_{h_1}(\omega)$ and $\omega_2 \in \mathcal{B}_{h_2}(\omega)$,

3. decision histories $h', h'' \in [h]$,

such that, for some $h'_1 \approx_i h_1$, $h'_1 \prec h' \prec \zeta(\omega_1)$ and, for some $h'_2 \approx_j h_2$, $h'_2 \prec h'' \prec \zeta(\omega_2)$.

That is, player i at h_1 considers it possible that the play has reached history $h'_1 \in [h_1]$ and, after taking an action at h'_1, information set $[h]$ of player k is reached, and player j at h_2 considers it possible that the play has reached history $h'_2 \in [h_2]$ and, after taking an action at h'_2, that same information set $[h]$ of player k is reached.

Definition 4.3. We say that at state ω active players i and j *are in agreement* if, for every future information set $[h]$ that they consider (Definition 4.2), they predict the same choices(s) of player $k = \iota(h)$ at h, that is, if player i is active at reached history h_1 and player j is active at reached history h_2, with $h_1 \prec h_2 \prec \zeta(\omega)$, then

1. if $\omega_1 \in \mathcal{B}_{h_1}(\omega)$ and $h'_1 \prec h'a \precsim \zeta(\omega_1)$ with $h'_1 \in [h_1]$, $h' \in [h]$ and $a \in A(h)$, then there exists an $\omega_2 \in \mathcal{B}_{h_2}(\omega)$ such that, for some $h'' \in [h]$ and $h'_2 \in [h_2]$, $h'_2 \prec h''a \precsim \zeta(\omega_2)$, and

2. if $\omega_2 \in \mathcal{B}_{h_2}(\omega)$ with $h'_2 \prec h''b \precsim \zeta(\omega_2)$ with $h'_2 \in [h_2]$, $h'' \in [h]$ and $b \in A(h)$ then here exists an $\omega_1 \in \mathcal{B}_{h_1}(\omega)$ such that, for some $h' \in [h]$ and $h'_1 \in [h_1]$, $h'_1 \prec h'b \precsim \zeta(\omega_1)$.

Finally we define the event, denoted by $\mathbf{A}$, that any two *active* players are in agreement:

$$\omega \in \mathbf{A} \text{ if and only if any two players active at } \omega \text{ are in agreement at } \omega. \tag{5}$$

Proposition 2. *Consider a finite extensive-form game G where no player moves more than once along any play. Then,*

(A) *If z is a Nash play of G then there is a causally restricted model of G and a state ω in that model such that (1) $\zeta(\omega) = z$ and (2) $\omega \in \mathbf{R} \cap \mathbf{T} \cap \mathbf{C} \cap \mathbf{A}$.*

(B) *For any causally restricted model of G and for every state ω in that model, if $\omega \in \mathbf{R} \cap \mathbf{T} \cap \mathbf{C} \cap \mathbf{A}$ then $\zeta(\omega)$ is a Nash play.*

The proof of Proposition 2 is given in the Appendix.

Remark 2. *Note that, in a perfect-information game, $\mathbf{T} \cap \mathbf{C} \subseteq \mathbf{A}$. Hence Proposition 1 is a corollary of Proposition 2.*

5 Conclusion

The characterization of Nash play given in Proposition 2, unlike characterizations of Nash equilibrium provided for strategic-form games, does not require players to believe in each other's rationality.[32] This can be seen in the game and model shown in Figure 9, where $\mathbf{R} = \{\gamma\}$, $\mathbf{T} = \{\beta, \gamma\}$ and $\mathbf{C} = \mathbf{A} = \{\alpha, \beta, \gamma\}$, so that $\mathbf{R} \cap \mathbf{T} \cap \mathbf{C} \cap \mathbf{A} = \{\gamma\}$ but at γ Player 1 does not believe that Player 2 is rational, because $\beta \in \mathcal{B}_{\emptyset}(\gamma)$ and at β Player 2 is not rational (she plays d believing that c gives her higher utility).

In the previous section we focused on the conditions that are needed to ensure that rational play in an extensive-form game coincides with Nash play. As noted above, these conditions do not require players to believe in each

[32] In their seminal paper Aumann and Brandenburger (1995) showed that, in games with more than two players, if there exists a common prior then mutual belief in rationality and payoffs as well as common belief in each player's conjecture about the opponents' strategies imply Nash equilibrium. However, Polak (1999) later showed that, in complete-information games, Aumann and Brandenburger's conditions actually imply common belief in rationality. Barelli (2009) generalized Aumann and Brandenburger's result by substituting the common prior assumption with the weaker property of action-consistency, and by replacing common belief in conjectures with a weaker condition stating that conjectures are constant in the support of the action-consistent distribution. He thus provided sufficient epistemic conditions for Nash equilibrium without requiring common belief in rationality. Later, Bach and Tsakas (2014) obtained a further generalization by introducing even weaker epistemic conditions for Nash equilibrium than those in Barelli (2009): their characterization of Nash equilibrium is based on introducing pairwise epistemic conditions imposed only on *some* pairs of players (contrary to the characterizations in Aumann and Brandenburger (1995) and Barelli (2009), which correspond to pairwise epistemic conditions imposed on *all* pairs of players). Not only do these conditions not imply common belief in rationality but they do not even imply "global" mutual belief in rationality; however the sufficient conditions provided by Bach and Tsakas (2014) still require pairwise mutual belief in rationality. Perea (2007) provides a one-person characterization of Nash strategies. For the case of perfect-information games, a comparison between the characterization of Nash play given in Proposition 1 and previous results in the literature is provided in (Bonanno 2018, Section 6).

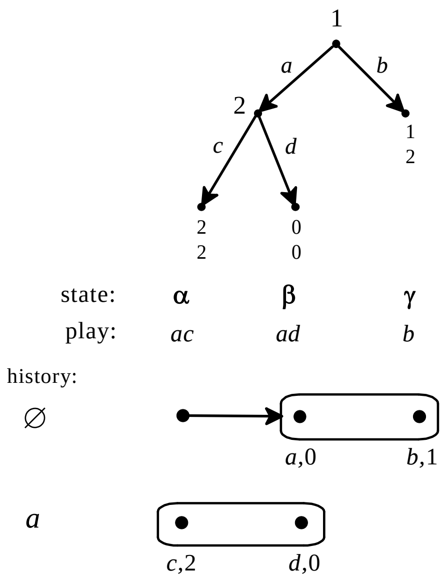

Figure 9: $\gamma \in \mathbf{R} \cap \mathbf{T} \cap \mathbf{C} \cap \mathbf{A}$ but at γ Player 1 does not believe that Player 2 is rational.

other's rationality. Hence they are too weak to yield stronger solution concepts, such as backward induction. Within the behavioral approach described in this paper, the conditions needed to obtain a characterization of backward induction in perfect information games, or a generalized version of it for games with imperfect-information, are investigated in Bonanno (2018) and Bonanno (2014), respectively. These conditions, which can be expressed as events, capture iterated belief in the rationality of future players, in the following sense: each active player is rational and believes that, after any of her actions, the players who will follow will be rational and will believe that, after any of their actions, the players who will follow will be rational, and so on. This is what Perea (2014) calls "belief in the opponents' future rationality" (a notion that he develops in the approach commonly used in the literature, where the underlying space of uncertainty is the set of the opponents' strategies) and Bonanno (2014) calls "forward belief in rationality" (developed within a dynamic behavioral approach, with a focus on von Neumann games, that is, games where any two histories that belong to the same information set have the same number of predecessors).

In Definition 3.4 we put forward the notion of self-confirming play, which is in the spirit of self-confirming equilibrium (Fudenberg and Levine (1993)), but framed in behavioral terms and without making use of the notion of strategy. We showed that in perfect-information games the notion of self-confirming play is equivalent to the notion of Nash play, but the equivalence does not extend to imperfect information games. Proposition 2 identified the additional restrictions that are needed to characterize the set of Nash plays in games with imperfect information.

A Proof of Proposition 2

Given a finite extensive-form game and a pure-strategy profile s, define the function $f_s : H \to Z$ (recall that H is the set of all histories and Z is the set of terminal histories) as follows: if $z \in Z$ then $f_s(z) = z$ and if $h \in D$ (recall that D is the set of decision histories) then $f_s(h)$ is the terminal history reached from h by following the choices prescribed by s. We denote by z_s^* the play generated by s, that is, the terminal history reached by s from the null history: $z_s^* = f_s(\emptyset)$. We say that z_s^* *avoids information set* $[h]$ if, for all $h' \in [h]$, $h' \nprec z_s^*$. If z_s^* does not avoid information set $[h]$ then we denote the unique history in $[h]$ that is a prefix of z_s^* by $h_s^*([h])$ (thus $h_s^*([h]) \in [h]$ and $h_s^*([h]) \prec z_s^*$).

Definition A.1. Given an extensive-form game G, denote by I the set of information sets. Let s be a pure-strategy profile of G. A *selection function based on* s is a function $g_s : I \to D$ that selects for every information set $[h] \in I$ a unique decision history in $[h]$ subject to the constraint that if z_s^* does not avoid information set $[h]$ then $g_s([h]) = h_s^*([h])$.

Definition A.2. Let G be an extensive-form game, s a pure strategy profile and g_s a selection function based on s. The *model of* G *generated by* s *and* g_s is the following model.

- $\Omega = Z$.
- $\zeta : Z \to Z$ is the identity function: $\zeta(z) = z, \forall z \in Z$.
- For every $h \in D$ and $z \in Z$ define $\mathcal{B}_h(z)$ as follows:
    1. If $h \nprec z$, then $\mathcal{B}_h(z) = \emptyset$.

2. If $h \prec z_s^*$ then $\mathcal{B}_h(z_s^*) = \{z' \in Z : z' = f_s(ha)$ for some $a \in A(h)\}$. [That is, if h is on the play generated by s, then at h the active player believes that, for every available action a, if she takes action a then the outcome will be the terminal history reached from ha by s.]
3. If $h \nprec z_s^*$, but $[h]$ is not avoided by z_s^*, then, for all $z \in Z$ such that $h \prec z$, $\mathcal{B}_h(z) = \{z' \in Z : z' = f_s(h_s^*([h])a)$ for some $a \in A(h)\}$.[That is, at every decision history in an information set crossed by the play generated by s, the player believes that the play has reached history $h_s^*([h])$ (the history in $[h]$ that is on the play to z_s^*) and her beliefs are as given in Point 2.]
4. If $[h]$ is avoided by z_s^*, let $\hat{h} = g_s([h])$. Then, for every $h' \in [h]$ and every $z \in Z$ such that $h' \prec z$, $\mathcal{B}_{h'}(z) = \{z' \in Z : z' = f_s(\hat{h}a)$ for some $a \in A(h)\}$. [That is, at every decision history in an information set that is not crossed by the play generated by s, the player believes that she is at the history selected by g_s, denoted by $\hat{h}$, and that, for every available action a, if she takes action a then the outcome will be the terminal history reached from $\hat{h}a$ by s.]

Remark 3. *Note that the model generated by a pure-strategy profile s and a selection function g_s is a causally restricted model (Definition 4.1).*

Remark 4. *Let G be a finite extensive-form game and consider the model generated by a pure-strategy profile s of G and a selection function g_s (Definition A.2). Then the no-uncertainty conditions 1 and 2 of Definition 3.3 and the agreement condition (5) are satisfied at every state, that is, $\mathbf{C} = \mathbf{A} = Z$. Furthermore, by Point 1 in Definition A.2, $z_s^* \in \mathcal{B}_h(z_s^*)$ for all h such that $h \prec z_s^*$; that is, $z_s^* \in \mathbf{T}$.*

We can now prove Proposition 2.

Proof. **(A)**[33] Fix a finite extensive-form game G and let s be a pure-strategy Nash equilibrium s of G. Fix a selection function g_s based on s (Definition A.1) and consider the model generated by s and g_s (Definition A.2). By Remark 4, $z_s^* \in \mathbf{C} \cap \mathbf{T} \cap \mathbf{A}$ (recall that z_s^* is the play generated by s, that is, $z_s^* = f_s(\emptyset)$). Thus it only remains to show that $z_s^* \in \mathbf{R}$. If h is a decision history, denote by $s(h)$ the choice selected by s at h. Fix an arbitrary decision history h that is reached at state $h \prec z_s^*$ (that is, $h \prec z_s^*$) and let a be the action at h such that $ha \precsim z_s^*$, that is, $s(h) = a$; then $f_s(ha) = f_s(\emptyset) = z_s^*$. Suppose that player $\iota(h)$ is not rational

[33]For this part of the proof we do not need the restriction that no player moves more than once along any play of the game.

at h. Then there is an action $b \in A(h) \setminus \{a\}$ that guarantees a higher utility to player $\iota(h)$, that is, if $z' \in \mathcal{B}_h(z_s^*)$ is such that $hb \precsim z'$, then $u_{\iota(h)}(z') > u_{\iota(h)}(z_s^*)$. By Definition A.2, $z' = f_s(hb)$ and thus $u_{\iota(h)}(f_s(hb)) > u_{\iota(h)}(f_s(ha))$ so that by unilaterally changing her strategy at h from a to b (and leaving the rest of her strategy unchanged), player $\iota(h)$ can increase her payoff, contradicting the assumption that s is a Nash equilibrium.

(B) Let G be a finite extensive-form game where no player moves more than once along any play and consider a model of it where there is a state α such that $\alpha \in \mathbf{R} \cap \mathbf{T} \cap \mathbf{C} \cap \mathbf{A}$. We need to construct a pure-strategy Nash equilibrium s of G such that $f_s(\emptyset) = \zeta(\alpha)$.

STEP 1. For every decision history h such that $h \prec \zeta(\alpha)$, let $s(h) = a$ where $a \in A(h)$ is the action at h such that $ha \precsim \zeta(\alpha)$.

STEP 2. Fix an arbitrary decision history h that is reached at state α (that is, $h \prec \zeta(\alpha)$) and an arbitrary $b \in A(h)$ such that $hb \not\precsim \zeta(\alpha)$ (that is, $b \neq s(h)$ where $s(h)$ was defined in Step 1). By Definition of model (Definition 3.1) there exists an $\hat{\omega} \in \mathcal{B}_h(\alpha)$ such that $\hat{h}b \precsim \zeta(\hat{\omega})$ for some $\hat{h} \in [h]$. Since $\alpha \in \mathbf{C}$, by Point 1 of Definition 3.3 for every $\omega' \in \mathcal{B}_h(\alpha)$ and for every $h' \in [h]$, if $h'b \precsim \zeta(\omega')$ then $h' = \hat{h}$. Since $\alpha \in \mathbf{C}$, by Point 2 of Definition 3.3 for any other $\omega \in \mathcal{B}_h(\alpha)$ such that $\hat{h}b \precsim \zeta(\omega)$, $\zeta(\omega) = \zeta(\hat{\omega})$. Define, for every h' such that $\hat{h}b \precsim h' \prec \zeta(\hat{\omega})$, $s(h') = c$ where $c \in A(h')$ is the action at h' such that $h'c \precsim \zeta(\hat{\omega})$. Note that, since $\alpha \in \mathbf{A}$, if any other active player at any reached history at state α considers the information set that contains history h', then that player will also predict choice c at h'. Thus $s(h')$ is well defined.

So far we have defined the choices prescribed by s along the play $\zeta(\alpha)$ and for paths to terminal histories following one-step deviations from this play.

STEP 3. Complete s in an arbitrary way.

Because of Step 1, $\zeta(\alpha) = f_s(h)$, for every $h \precsim \zeta(\alpha)$ (in particular, $f_s(\emptyset) = \zeta(\alpha)$). We want to show that s is a Nash equilibrium. Suppose not. Then there is a decision history h with $h \prec \zeta(\alpha)$ such that, by changing her choice at h from $s(h)$ to a different choice, player $\iota(h)$ can increase her payoff (recall the assumption that the game satisfies the no-consecutive-moves assumption and thus there are no successors of h that belong to player $\iota(h)$). Let $s(h) = a$ (thus $ha \precsim \zeta(\alpha)$) and let b be the choice at h that yields a higher payoff to player $\iota(h)$; that is,

$$u_{\iota(h)}(f_s(hb)) > u_{\iota(h)}(\zeta(\alpha)). \tag{6}$$

Let $\omega \in \mathcal{B}_h(\alpha)$ be such that $hb \precsim \zeta(\omega)$ (such an ω exists by Point 4 of Definition

3.1). Since $\alpha \in \mathbf{C}$, for every $\omega' \in \mathcal{B}_h(\alpha)$ such that $hb \precsim \zeta(\omega')$, $\zeta(\omega) = \zeta(\omega')$. By Step 2 above,

$$\zeta(\omega) = f_s(hb). \tag{7}$$

It follows from (7) that, at state α and history h, player $\iota(h)$ believes that if she plays b her payoff will be $u_{\iota(h)}(f_s(hb))$. Since $\alpha \in \mathbf{T}$, $\alpha \in \mathcal{B}_h(\alpha)$, and since $\alpha \in \mathbf{C}$, for every $\omega' \in \mathcal{B}_h(\alpha)$ such that $ha \precsim \zeta(\omega')$, $\zeta(\omega') = \zeta(\alpha)$. Thus, at state α and history h, player $\iota(h)$ believes that if she plays a her payoff will be $u_{\iota(h)}(\zeta(\alpha))$. It follows from this and (6) that at α and h player $\iota(h)$ believes that action b is better than action a, which implies that at α player $\iota(h)$ is not rational, contradicting the assumption that at $\alpha \in \mathbf{R}$. □

References

R. Aumann. Correlated equilibrium as an expression of Bayesian rationality. *Econometrica*, 55:1–18, 1987.

R. Aumann. Backward induction and common knowledge of rationality. *Games and Economic Behavior*, 8:6–19, 1995.

R. Aumann. On the centipede game. *Games and Economic Behavior*, 23:97–105, 1998.

R. Aumann and A. Brandenburger. Epistemic conditions for Nash equilibrium. *Econometrica*, 63:1161–1180, 1995.

C. Bach and E. Tsakas. Pairwise epistemic conditions for Nash equilibrium. *Games and Economic Behavior*, 85:48–59, 2014.

A. Baltag, S. Smets, and J. Zvesper. Keep 'hoping' for rationality: a solution to the backward induction paradox. *Synthese*, 169:301–333, 2009.

P. Barelli. Consistency of beliefs and epistemic conditions for Nash and correlated equilibria. *Games and Economic Behavior*, 67:363–375, 2009.

P. Battigalli. Comportamento razionale ed equilibrio nei giochi e nelle situazioni strategiche. Unpublished dissertation, Bocconi University, Milano, 1987.

P. Battigalli and G. Bonanno. Recent results on belief, knowledge and the epistemic foundations of game theory. *Research in Economics*, 53:149–225, 1999.

P. Battigalli and D. Guaitoli. Conjectural equilibria and rationalizability in a game with incomplete information. In P. Battigalli, A. Montesano, and F. Panunzi, editors, *Decisions, Games and Markets*, pages 97–124. Kluwer Academic, 1997.

P. Battigalli and M. Siniscalchi. Strong belief and forward induction reasoning. *Journal of Economic Theory*, 106:356–391, 2002.

P. Battigalli, A. Di-Tillio, and D. Samet. Strategies and interactive beliefs in dynamic games. In D. Acemoglu, M. Arellano, and E. Dekel, editors, *Advances in Economics and Econometrics. Theory and Applications: Tenth World Congress, Volume 1*, pages 391–422. Cambridge University Press, Cambridge, 2013.

D. Bernheim. Rationalizable strategic behavior. *Econometrica*, 52:1002–1028, 1984.

C. Bicchieri and M. Green. Symmetry arguments for cooperation in the Prisoner's Dilemma. In C. Bicchieri, R. Jeffrey, and B. Skyrms, editors, *The logic of strategy*, pages 175–195. Oxford University Press, Oxford, 1999.

G. Bonanno. A dynamic epistemic characterization of backward induction without counterfactuals. *Games and Economics Behavior*, 78:31–43, 2013a.

G. Bonanno. *AGM*-consistency and perfect Bayesian equilibrium. Part I: definition and properties. *International Journal of Game Theory*, 42:567–592, 2013b.

G. Bonanno. A doxastic behavioral characterization of generalized backward induction. *Games and Economics Behavior*, 88:221–241, 2014.

G. Bonanno. Epistemic foundations of game theory. In H. van Ditmarsch, J. Halpern, W. van der Hoek, and B. Kooi, editors, *Handbook of Epistemic Logic*, pages 443–488. College Publications, 2015a.

G. Bonanno. Reasoning about strategies and rational play in dynamic games. In J. van Benthem, S. Ghosh, and R. Verbrugge, editors, *Modeling strategic reasoning*, pages 34–62. Springer, 2015b.

G. Bonanno. AGM-consistency and perfect Bayesian equilibrium. Part II: from PBE to sequential equilibrium. *International Journal of Game Theory*, 45:1071–1094, 2016.

G. Bonanno. Behavior and deliberation in perfect-information games: Nash equilibrium and backward induction. *International Journal of Game Theory*, 47: 1001–1032, 2018.

G. Bonanno. The material conditional is sufficient to model deliberation. *Erkenntnis*, 2021. doi: 10.1007/s10670-020-00357-7.

A. Brandenburger and E. Dekel. Rationalizability and correlated equilibria. *Econometrica*, 55:1391–1402, 1987.

E. Dekel, D. Fudenberg, and D. Levine. Payoff information and self-confirming equilibrium. *Journal of Economic Theory*, 89:165–185, 1999.

E. Dekel, D. Fudenberg, and D. Levine. Subjective uncertainty over behavior strategies: a correction. *Journal of Economic Theory*, 104:473–478, 2002.

I. Esponda. Rationalizable conjectural equilibrium: A framework for robust predictions. *Theoretical Economics*, 8:467–501, 2013.

D. Fudenberg and Y. Kamada. Rationalizable partition-confirmed equilibrium. *Theoretical Economics*, 10:775–806, 2015.

D. Fudenberg and D. Levine. Self-confirming equilibrium. *Econometrica*, 61: 523–545, 1993.

D. Fudenberg and J. Tirole. Perfect Bayesian equilibrium and sequential equilibrium. *Journal of Economic Theory*, 53:236–260, 1991.

D. Gauthier. *Morals by agreement*. Oxford University Press, 1986.

M. Gilli. Metodo bayesiano e aspettative nella teoria dei giochi e nella teoria economica. Unpublished dissertation, Bocconi University, Milano, 1987.

M. Gilli. On non-Nash equilibria. *Games and Economic Behavior*, 27:184–203, 1999.

J. Greenberg. The right to remain silent. *Theory and Decision*, 48:193–204, 2000.

J. Greenberg, S. Gupta, and X. Luo. Mutually acceptable courses of action. *Economic Theory*, 40:91–112, 2009.

J. Halpern. Hypothetical knowledge and counterfactual reasoning. *International Journal of Game Theory*, 28:315–330, 1999.

J. Halpern. Substantive rationality and backward induction. *Games and Economic Behavior*, 37:425–435, 2001.

D. Kreps and R. Wilson. Sequential equilibrium. *Econometrica*, 50:863–894, 1982.

D. Lewis. *Counterfactuals*. Harvard University Press, 1973.

R. Myerson. *Game Theory*. Harvard University Press, 1991.

R. Nozick. Newcomb's problem and two principles of choice. In N. Rescher, editor, *Essays in Honor of Carl G. Hempel: A Tribute on the Occasion of his Sixty-Fifth Birthday*, pages 114–146. Springer Netherlands, Dordrecht, 1969. doi: 10.1007/978-94-017-1466-2_7.

M. Osborne and A. Rubinstein. *A course in game theory*. MIT Press, Cambridge, 1994.

D. Pearce. Rationalizable strategic behavior and the problem of perfection. *Econometrica*, 52:1029–1050, 1984.

A. Penta. Robust dynamic mechanism design. Technical report, University of Wisconsin, Madison, 2009.

A. Perea. A one-person doxastic characterization of Nash strategies. *Synthese*, 158:251–271, 2007.

A. Perea. Backward induction *versus* forward induction reasoning. *Games*, 1: 168–188, 2010.

A. Perea. *Epistemic game theory: reasoning and choice*. Cambridge University Press, Cambridge, 2012.

A. Perea. Belief in the opponents' future rationality. *Games and Economic Behavior*, 83:231–254, 2014.

B. Polak. Epistemic conditions for Nash equilibrium, and common knowledge of rationality. *Econometrica*, 67:673–676, 1999.

A. Rubinstein and A. Wolinsky. Rationalizable conjectural equilibrium: between Nash and rationalizability. *Games and Economic Behavior*, 6:299–311, 1994.

D. Samet. Hypothetical knowledge and games with perfect information. *Games and Economic Behavior*, 17:230–251, 1996.

W. Spohn. Dependency equilibria and the causal structure of decision and game situations. *Homo Oeconomicus*, 20:195–255, 2003. doi: 10.1007/s41412-017-0060-5.

W. Spohn. Dependency equilibria. *Philosophy of Science*, 74(5):775–789, 2007.

W. Spohn. From nash to dependency equilibria. In G. Bonanno, B. Löwe, and W. van der Hoek, editors, *Logic and the Foundations of Game and Decision Theory – LOFT 8*, pages 135–150, Berlin, Heidelberg, 2010. Springer.

R. Stalnaker. A theory of conditionals. In N. Rescher, editor, *Studies in logical theory*, pages 98–112. Blackwell, 1968.

R. Stalnaker. On the evaluation of solution concepts. *Theory and Decision*, 37: 49–74, 1994.

R. Stalnaker. Knowledge, belief and counterfactual reasoning in games. *Economics and Philosophy*, 12:133–163, 1996.

R. Stalnaker. Belief revision in games: forward and backward induction. *Mathematical Social Sciences*, 36:31–56, 1998.

R. Stalnaker and R. Thomason. A semantical analysis of conditional logic. *Theoria*, 36:246–281, 1970.

T. Tan and S. Werlang. The Bayesian foundation of solution concepts of games. *Journal of Economic Theory*, 45:370–391, 1988.