

Article

Rational Play in Extensive-Form Games

Giacomo Bonanno 

Department of Economics, University of California, Davis, CA 95616-8578, USA; gfbonanno@ucdavis.edu

Abstract: We argue in favor of a departure from the equilibrium approach in game theory towards the less ambitious goal of describing only the actual behavior of rational players. The notions of Nash equilibrium and its refinements require a specification of the players' choices and beliefs not only along the equilibrium play but also at counterfactual histories. We discuss an alternative—counterfactual-free—approach that focuses on choices and beliefs along the actual play, while being silent on choices and beliefs at unreached histories. Such an approach was introduced in an earlier paper that considered only perfect-information games. Here we extend the analysis to general extensive-form games (allowing for imperfect information) and put forward a behavioral notion of self-confirming play, which is close in spirit to the literature on self-confirming equilibrium. We also extend, to general extensive-form games, the characterization of rational play that is compatible with pure-strategy Nash equilibrium.

Keywords: material rationality; behavioral model; self-confirming play; Nash equilibrium

1. Introduction

We address the issue of what kind of object qualifies as a “rational solution” of an extensive-form game. Whereas the dominant approach focuses on strategy profiles that, besides being Nash equilibria, satisfy additional—often increasingly complex—criteria, we suggest moving in the opposite direction by doing without the notion of strategy and focusing only on the actions and beliefs of the active players, that is, on choices made and beliefs held at reached histories.

The notions of Nash equilibrium and its refinements require a specification of the players' choices and beliefs not only along the equilibrium play but also at counterfactual (that is, unreached) histories. In Section 2, we argue that pinning down counterfactual choices at unreached information histories is not a straightforward matter and that it may be futile to search for a general theory of rationality that can achieve this result. Instead, we argue in favor of a more basic approach that was put forward in [1]; while [1] was exclusively focused on perfect-information games, in this paper we extend the approach to general extensive-form games, by allowing for imperfect information. We put forward a behavior-based notion of self-confirming play—which is close in spirit to the literature on self-confirming equilibrium—and extend the characterization of rational play that is compatible with pure-strategy Nash equilibrium to general extensive-form games.

The paper is organized as follows. In Section 2, we illustrate, by means of examples, the difficulties that arise in the pursuit of specifying rational choices and beliefs at unreached information sets. In Section 3, we review the behavioral models introduced in [1], extend them to general extensive-form games and put forward a notion of self-confirming play, where each action taken is justified by the beliefs held at the time of choice and, furthermore, those beliefs turn out to be exactly correct, so that no player receives information that contradicts those beliefs (and thus experiences no regret). We also extend the characterization of rational play that is consistent with pure-strategy Nash equilibrium to general extensive-form games. Section 4 provides further discussion and a conclusion.



Citation: Bonanno, G. Rational Play in Extensive-Form Games. *Games* **2022**, *13*, 72. <https://doi.org/10.3390/g13060072>

Academic Editor: Ulrich Berger

Received: 24 September 2022

Accepted: 26 October 2022

Published: 30 October 2022

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2022 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

2. What Is a Rational Solution?

What constitutes a rational solution of an extensive-form game? Two different approaches can be found in the literature.

1. The Nash-equilibrium approach. Within this approach the notion of rationality is captured through the concept of Nash equilibrium or one of its refinements. Consider, for example, the game of Figure 1 and the strategy profile (a, a, d) .

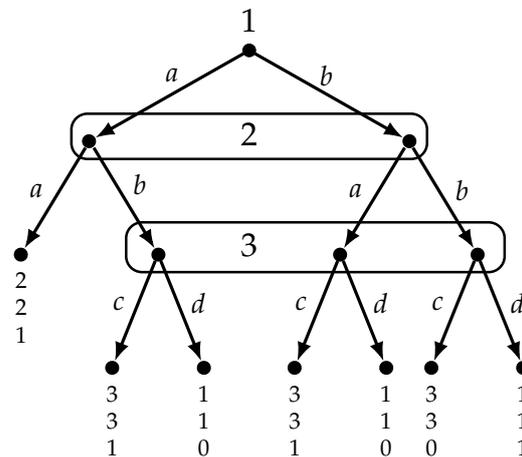


Figure 1. An extensive game with imperfect information.

Although (a, a, d) is a Nash equilibrium, popular refinements of Nash equilibrium would deny it the status of a “rational solution” [for example, (a, a, d) is not a sequential equilibrium [2] because strategy d can be a rational choice for Player 3 only if she assigns positive probability to history bb , but the notion of consistency—which is part of the definition of sequential equilibrium—requires Player 3’s beliefs to assign zero probability to bb]. Regardless of one’s views on whether (a, a, d) can be considered a rational solution of the game of Figure 1, there is a more fundamental issue to be considered, namely how Player 3’s strategy d should be interpreted. The common interpretation seems to be in terms of an objective counterfactual: *Player 3 would play d if her information set were to be reached*. It is typically the case in extensive-form games that, given a strategy profile s , there will be information sets that are not reached by the play generated by s . Thus, under this interpretation of strategies, a rational solution of the game would determine, not only what actions are actually taken by the players (that is, what the actual play of the game is), but also—counterfactually—what actions would be taken at every unreached information set.

The standard theory of counterfactuals, due to Robert Stalnaker and David Lewis [3–5], 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. Referring to the game of Figure 1, at a world where Players 1 and 2 play aa , we can take ϕ to be the sentence “Player 3’s information set is reached” and ψ the sentence “Player 3 plays d ”. 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 aa) if and only if the most similar world to the actual world at which Player 3’s information is reached is one where Player 3 plays d . However, how are we to determine if the most similar world to the actual world is one where Player 3 plays d or one where Player 3 plays c ?

In general, pinning down counterfactual choices at unreached information sets is not a straightforward matter. Consider, for example, the game illustrated in Figure 2 due to Perea ([6], p. 169), where Player 1 can either play b and end the game, or play a , in which case Players 1 and 2 play a “simultaneous” game.

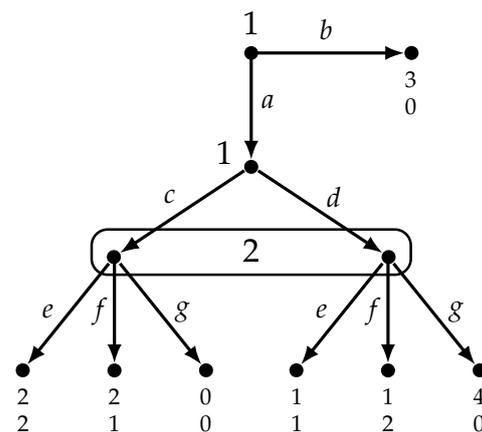


Figure 2. The conflict between the backward-induction-based counterfactual and the forward-induction-based counterfactual encoded in Player 2’s strategy.

If one appeals to backward-induction reasoning, one is led to conclude that the “rational solution” of this game is the strategy profile (b, e) , which incorporates the counterfactual claim that Player 2 would play e if her information set were to be reached [first apply the procedure of iterative deletion of strictly dominated strategies to the subgame that starts at history a to obtain (c, e) : 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 ; then infer that Player 1 will play b ; backward-induction reasoning is captured by such notions as “common belief in present and future rationality” [7], or forward belief in rationality [8,9].]

On the other hand, if one appeals to forward-induction reasoning (as captured by the notion of extensive-form rationalizability [10,11]) one is led to conclude that the “rational solution” of this game is the strategy profile (b, f) , incorporating the counterfactual claim that Player 2 would play f if her information set were to be reached [first eliminate Player 1’s strategy ac , since it is strictly dominated by b , and Player 2’s strategy g ; then eliminate Player 1’s strategy ad and Player 2’s strategy e , with the conclusion that Player 1 will play b and Player 2 would play f .]

Note, however, that the prediction in terms of play is the same, namely that Player 1 would end the game by playing b ; the two solutions differ only in terms of the answer to the question “what would Player 2 do if her information set were to be reached?”

The above example shows that it may be futile to search for a general theory of rationality that would pin down counterfactual choices and beliefs at unreached information sets (as shown recently by [12], besides backward-induction and forward-induction reasoning, there are other types of rationality-based reasoning that lead to the same outcome but different counterfactual “predictions” about choices at unreached information sets). It is natural, therefore, to ask: Is it essential to provide an answer to such counterfactual questions? The thesis put forward in this paper is that the answer to this question is negative.

2. The self-confirming equilibrium approach. Returning to the strategy profile (a, a, d) in the game of Figure 1, an alternative approach is to interpret Player 3’s strategy d not as a claim about what Player 3 would actually do in a counterfactual world where her information is reached, but as a belief, shared by Players 1 and 2, about Player 3’s hypothetical behavior. Such shared belief would support the rationality of playing a for both Players 1 and 2.

This approach is in line with the literature that identifies rational play in extensive-form games with the notion of self-confirming equilibrium, introduced in [13] (similar notions are put forward in [14–16]; Refs. [17,18] 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 [19] provide a generalization of self-confirming equilibrium; the related expression ‘conjunctural equilibrium’ is mostly used in the context

of strategic-form games: it was introduced in this context by [20,21] and is defined as a situation where each player’s strategy is a best response to a conjecture about the other players’ strategies and any information acquired after the play of the game does not induce the player to change her conjecture). A self-confirming equilibrium is a strategy profile satisfying the property that 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, even though beliefs about play at unreached information sets may be incorrect. The essential feature of a self-confirming equilibrium is that no player receives information that contradicts her beliefs.

If one follows the interpretation suggested above, then two issues arise. First of all, the strategy profile (a, a, d) (for the game of Figure 1) is now a hybrid object, incorporating—on the one hand—a prediction about actual behavior (namely, the (a, a) part) and—on the other hand – an encoding of the beliefs of Players 1 and 2 (namely, the d part). This leaves to be desired, since one should clearly distinguish between actions and beliefs and model the latter explicitly. The second issue is that there seems to be no reason to require different players to agree on the hypothetical choice of a third player at an unreached information set. Consider, for example, the game of Figure 3, taken from ([13], p. 533).

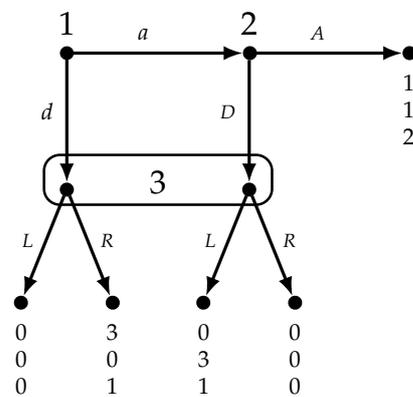


Figure 3. The play aA is consistent with the notion of self-confirming equilibrium, even though there is no Nash equilibrium that yields aA .

In this game it is rational for Player 1 to play a , if she believes that Player 2 will play A and Player 3 would play L , and it is rational for Player 2 to play A , if he believes that Player 3 would play R . Thus, the play aA is supported by beliefs of Players 1 and 2 that are not in agreement with each other.

Note, however, that the notion of self-confirming equilibrium is still defined in terms of a strategy profile and thus one cannot claim that (a, A) is a self-confirming equilibrium. One would have to state that 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 would play R . In other words, also the notion of self-confirming equilibrium requires an answer to the counterfactual “what would Player 3 do if her information set were to be reached?” Note also that in the game of Figure 3 there is no Nash equilibrium that yields the play aA ; thus, a self-confirming equilibrium need not be a Nash equilibrium.

Both the notion of Nash equilibrium and the notion of self-confirming equilibrium require specifying choices at all information sets, whether they are reached or not. 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 3 and a situation where Player 1 plays a – believing that Player 2 will play A and 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 and thus which of Players 1 and 2 is holding incorrect beliefs? [Note that,

as [15] points out, in this game 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 .] 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.

In this paper, we turn to an alternative approach, put forward in [1] in the context of perfect-information games, and extend it to general extensive-form games. The proposed framework restricts attention to the *actual* choices of the players and the beliefs that justify those choices.

3. Behavioral Models of Games

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 (note that there is a straightforward way of translating one type of model into the other). 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, thus providing an interpretation of a state in terms of players’ choices. If ω is a state and s_i is the strategy of player i at ω 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 Players 1 and 2 play aa , so that Player 3’s information set is not reached; suppose also that the strategy of Player 3 associated with state ω is d . In what sense does Player 3 “play” d ? Does it mean that, before the game starts, Player 3 has made a plan to play d if her information happens 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 d ? [This interpretation is adopted in [22] where it is pointed 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”.] Or is Player 3’s strategy d to be interpreted not as a statement about what Player 3 would do but as an expression of what the opponents think that Player 3 would do?

While most of the literature on the epistemic foundations of game theory makes use of strategy-based models, a few papers follow a *behavioral* approach by associating with each state a *play* (or outcome) of the game (the seminal contribution is [23], followed by [1,8,24,25]; the focus of this literature has been on games with perfect information). 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 ([23] uses extended information structures to model hypothetical knowledge, [8] use plausibility relations and [24] use conditional probability systems). Here we follow the simpler approach put forward in [1], which models the “pre-choice” beliefs of a player, while the previous literature considered the “after-choice” beliefs. The previous literature was based on the assumption that, if at a state a player takes action a , then she knows that she takes action a , that is, in all the states that she considers possible she takes action a . The pre-choice or deliberation stage approach, on the other hand, models the beliefs of the player at the time when she is contemplating the actions available to her and treats each of those actions as an “open possibility”. Thus, her beliefs take the 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 ”, where the conditional “if p then q ” is interpreted as a *material conditional*, that is, as equivalent to “either not p or q ” (in [26] it is argued that, contrary to a common view, the material conditional is indeed sufficient to model deliberation; it is also shown how to convert pre-choice beliefs into after-choice beliefs, reflecting a later stage at which the agent has made up her mind on what to do). This analysis does not rely in any way on

counterfactuals; furthermore, only the beliefs of the active players at the time of choice are modeled, so that *no initial beliefs nor belief revision policies are postulated*. The approach is described below and it makes use of the history-based definition of extensive-form game, which is reviewed in Appendix A.

As in [1] we take a non-quantitative approach based on qualitative beliefs and ordinal utility.

3.1. Qualitative Beliefs

Let Ω be a set, whose elements are called *states* (or possible worlds). We represent the beliefs of an agent by means of 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 $\mathcal{B}(\omega) \subseteq \mathcal{B}(\omega')$ (it is well known that transitivity of \mathcal{B} corresponds to positive introspection of beliefs: if the agent believes an event \mathbf{E} then she believes that she believes \mathbf{E} , and euclideanness corresponds to negative introspection: if the agent does not believe \mathbf{E} then she believes that she does not believe \mathbf{E}). We will assume throughout that the belief relations are transitive and euclidean so that $\omega' \in \mathcal{B}(\omega)$ implies that $\mathcal{B}(\omega') = \mathcal{B}(\omega)$. Note that we do not assume reflexivity of \mathcal{B} (that is, we do not assume that, for every state ω , $\omega \in \mathcal{B}(\omega)$); reflexivity corresponds to the assumption that a player cannot have incorrect beliefs: an assumption that, as [27] points out, is conceptually problematic, especially in a multi-agent context). Hence, in general, the relation \mathcal{B} does *not* induce a partition of the set of states.

Graphically, we represent a transitive and euclidean belief relation as shown in Figure 4, where $\omega' \in \mathcal{B}(\omega)$ if and only if either there is an arrow from ω to the rounded rectangle containing ω' , or ω and ω' are enclosed in the same rounded rectangle (that is, if there is an arrow from state ω to a rounded rectangle, then, for every ω' in that rectangle, $(\omega, \omega') \in \mathcal{B}$ and, for any two states ω and ω' that are enclosed in a rounded rectangle, $\{(\omega, \omega), (\omega, \omega'), (\omega', \omega), (\omega', \omega')\} \subseteq \mathcal{B}$).

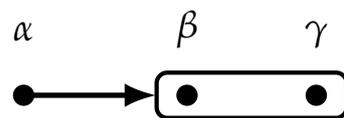


Figure 4. The relation $\mathcal{B} = \{(\alpha, \beta), (\alpha, \gamma), (\beta, \beta), (\beta, \gamma), (\gamma, \beta), (\gamma, \gamma)\}$.

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 $\{\beta, \gamma\}$. 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 $\{\beta, \gamma\}$, since event $\{\beta, \gamma\}$ is not true at α ($\alpha \notin \mathcal{B}(\alpha) = \{\beta, \gamma\}$). We say that *at state ω the agent has correct beliefs* if $\omega \in \mathcal{B}(\omega)$ (note that it is a consequence of euclideanness of the relation \mathcal{B} that, even if the agent’s beliefs are objectively incorrect, *she always believes that what she believes is true*: if $\omega' \in \mathcal{B}(\omega)$ then $\omega' \in \mathcal{B}(\omega')$).

3.2. Models of Games

As a starting point in the definition of a model of a game, we take a set of states Ω and provide an interpretation of each state in terms of a particular play of the game, by means of a function $\zeta : \Omega \rightarrow Z$ that associates, with every state ω , a play or terminal history $\zeta(\omega) \in Z$. Each state ω also provides a description of the beliefs of the active players by means of a binary relation \mathcal{B}_h on Ω representing the beliefs of $i(h)$, the player who moves at decision history h . It would be more precise to write $\mathcal{B}_{i(h)}$ instead of \mathcal{B}_h , but we have chosen the lighter notation since there is no ambiguity, because we assume (see Appendix A) that

at every decision history there is a unique player who is active there. Note that beliefs are specified only at histories that are *reached* at a given state, in the sense that $\mathcal{B}_h(\omega) \neq \emptyset$ if and only if $h \prec \zeta(\omega)$.

Definition 1. A model of an extensive-form game is a tuple $\langle \Omega, \zeta, \{\mathcal{B}_h\}_{h \in D} \rangle$ where

- Ω is a set of states.
- $\zeta : \Omega \rightarrow Z$ is an assignment of a terminal history to each state.
- For every $h \in D$, $\mathcal{B}_h \subseteq \Omega \times \Omega$ is a belief relation that satisfies the following properties:
 1. $\mathcal{B}_h(\omega) \neq \emptyset$ if and only if $h \prec \zeta(\omega)$ [beliefs are specified only at reached decision histories and are consistent: consistency means that there is no event \mathbf{E} such that both \mathbf{E} and its complement $\neg \mathbf{E}$ are believed; it is well known that, at state ω , beliefs are consistent if and only if $\mathcal{B}(\omega) \neq \emptyset$].
 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; recall (see Appendix A) 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 $\iota(h) = \iota(h')$)].
 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, which is what we do in Figures 5–9].
 4. If $\omega' \in \mathcal{B}_h(\omega)$ and $h' \prec \zeta(\omega')$ with $h' \approx_{\iota(h)} h$, then, for every action $a \in A(h)$ (note that $A(h') = A(h)$), there is an $\omega'' \in \mathcal{B}_h(\omega)$ such that $h'a \lesssim \zeta(\omega'')$.

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' ($h'a \lesssim \zeta(\omega'')$). This means that, for every available action, the active player at h considers it possible that she takes that action and thus has a belief about what will happen conditional on taking it. A further “natural” restriction on beliefs will be discussed later (Definition 6).

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

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

1. As a matter of fact, Player 1 plays a , Player 2 plays b and Player 3 plays d .
2. Player 1 (who chooses at the null 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 d and thus her utility will be 1.
3. Player 2 (who chooses 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 d and thus his utility will be 1, and believes that if he plays a his utility will be 2.
4. Player 3 (who chooses at information set $\{ab, ba, bb\}$) erroneously believes that both Player 1 and Player 2 played b ; thus, she believes that if she plays c her utility will be 0 and if she plays d her utility will be 1.

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 (that is, Player 3 is not active at state β since her information set is not reached).

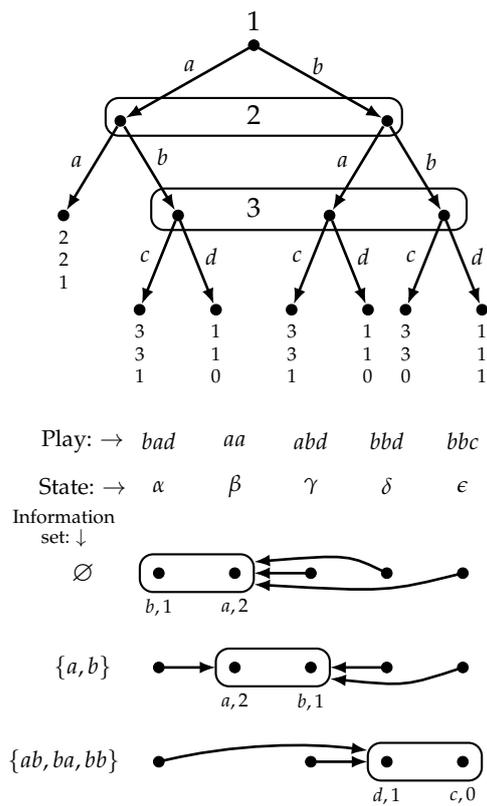


Figure 5. The top part reproduces the game of Figure 1 and the bottom part shows a model of it.

3.3. Rationality

Consider again the model of Figure 5 and state γ . There Player 1 believes that if she takes action a , her utility will be 2, and if she takes action b , her utility will be 1. 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 the outcome turns out to be abd not aa so that her utility will be 1, not 2). 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 those states where she takes action d , namely states α , γ and δ . 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* [28,29]. The notion of material rationality is the weaker of the two in that a player can be found to be irrational only at decision histories of hers that are actually reached (substantive rationality, on the other hand, is more demanding since a player can be labeled as irrational at a decision history h of hers even if h not reached). Given that we have adopted a purely behavioral approach, the natural notion for us is the weaker one, namely 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, $ha \succsim \zeta(\omega)$) then there is no other action at h that, according to her beliefs, *guarantees* a higher utility.

Definition 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 all h' such that $h' \approx_{\iota(h)} h$ (that is, history h' belongs to the same information set as h), if a is the action taken at history h' at state ω_1 , that is, $h'a \succsim \zeta(\omega_1)$, and b is the action taken at h' at state ω_2 , that is, $h'b \succsim \zeta(\omega_2)$, then $u_{\iota(h)}(\zeta(\omega_1)) < u_{\iota(h)}(\zeta(\omega_2))$. Thus, the active player at history h believes that action b is better than action a if, restricting attention to the states that she considers possible, the largest utility that she obtains if she plays a is less than the lowest 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 $ha \succsim \zeta(\omega)$ (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:

$$\omega \in \mathbf{R} \text{ if and only if, for every } h \prec \zeta(\omega), \text{ player } \iota(h) \text{ is rational at } h \text{ (at state } \omega). \tag{1}$$

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

3.4. Correct Beliefs

The notion of correct belief was first mentioned in Section 3.1 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 \mathbf{T} ('T' stands for 'true'), is defined as follows:

$$\omega \in \mathbf{T} \text{ if and only if } \omega \in \mathcal{B}_h(\omega) \text{ for every } h \text{ such that } h \prec \zeta(\omega). \tag{2}$$

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 (2) ($\beta \in \mathcal{B}_\emptyset(\beta)$ and $\beta \in \mathcal{B}_{\{a,b\}}(\beta)$). 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 , and (ii) she believes that if she plays b then Players 2 and 3 will play a and d , 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 *bad*" (which is equivalent to "either not b or *bad*") 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 d , 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 [3,4]); 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 d , respectively. Note that there is no a priori reason to think that the closest state to β where Player 1 plays b is state α . This is because, as pointed out by Stalnaker ([30], p.48), there is no necessary connection between counterfactuals, which capture causal relations, and beliefs: for example, I may 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 my belief is wrong because—as a matter of fact—if I were to drop it, it would not break, since it is 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 d , respectively?" One could, of course, enrich the models in order to answer the question, 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 states where action a is not taken are not relevant for the truth of the conditional.

3.5. Self-Confirming Play

We have defined two events: the event \mathbf{R} that all the *active* players are rational and the event \mathbf{T} that all the *active* players have 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 (namely, (a, a, d)) whose associated play is aa . 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 will 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 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, that is, there is no Nash equilibrium whose associated play is ad .

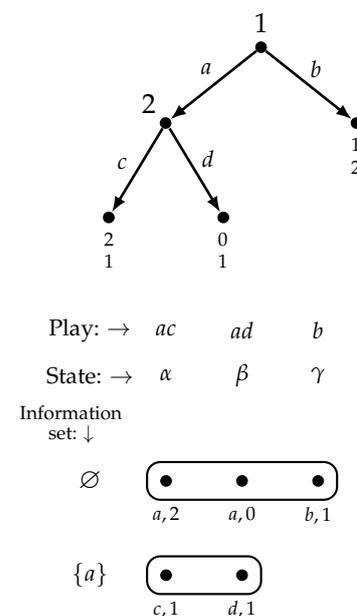


Figure 6. A perfect-information game and a model of it.

Next we introduce another event which, in conjunction with **T**, guarantees that the active players' beliefs about the opponents' actual moves are *exactly correct* (note that a requirement built in the definition of a self-fulfilling equilibrium ([13], p.523) is that "each player's beliefs about the opponents' play are exactly correct"). 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. A state ω belongs to event **C** if and only if, for every reached history h 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 $\zeta(\omega') = \zeta(\omega'')$.

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 $h' \neq h$ (for example, in the model of Figure 5, at state γ , which belongs to event **C**, and at 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), and—concerning Point 2—a player may also be erroneous in her certainty about what will happen after her choice (for example, in the model of Figure 5, at state γ and history \emptyset , 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).

In the model of Figure 5 $\mathbf{C} = \Omega$, while in the model of Figure 6 $\mathbf{C} = \emptyset$, because at the null history \emptyset Player 1 is uncertain about what will happen if she takes action a .

If state ω belongs to the intersection of events **C** and **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 3.4—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 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 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 ([1], Proposition 1, p. 1012). 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. 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 3 together with a model of it.

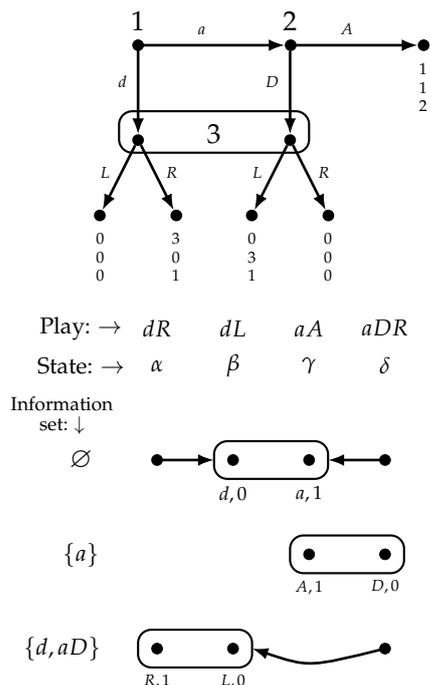


Figure 7. The game of Figure 3 and a model of it.

In the model of Figure 7, $\mathbf{R} = \{\gamma\}$, $\mathbf{T} = \{\beta, \gamma\}$ and $\mathbf{C} = \Omega$, so that $\mathbf{R} \cap \mathbf{T} \cap \mathbf{C} = \{\gamma\}$. Thus, at state γ the active players (Players 1 and 2) are rational, have correct beliefs and have no uncertainty and yet $\zeta(\gamma) = aA$ which is not a Nash play (there is no Nash equilibrium that yields the play aA). Players 1 and 2 have different beliefs about what Player 3 would do at her information set: at state γ Player 1 believes that if she plays d then Player 3 will play L , while Player 2 believes that if he plays D then Player 3 will play R .

Next we introduce a new event, denoted by \mathbf{A} (\mathbf{A} stands for ‘agreement’), that rules out such disagreement and use it to 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 (Definition 1) allows for “unreasonable” beliefs that express a causal link between a player’s action and her opponent’s reaction to it, when the latter does not observe the former’s choice. As an illustration of such beliefs, consider a game where Player 1 moves first, choosing between actions a and b , and Player 2 moves second choosing between actions 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 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 ought to be rejected as “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 and thus cannot react differently to Player 1’s choice of a and Player 1’s choice of b . [It should be noted, however, that several authors have argued that such beliefs are not necessarily irrational: see, for example, [31–36]]. 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 6. A causally restricted model of a game is a model (Definition 1) that satisfies the following additional restriction (a verbal interpretation follows; note that, for games with perfect information, there is no difference between a model and a restricted model, since (3) is vacuously satisfied).

- 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 $h' \approx_{\iota(h)} h$):

$$\begin{aligned} &\text{if } h', h'' \in [h], \omega_1, \omega_2 \in \mathcal{B}_h(\omega), h'a \prec h_1c_1 \lesssim \zeta(\omega_1) \text{ and} \\ &h''b \prec h_2c_2 \lesssim \zeta(\omega_2), \text{ then either } c_1 = c_2 \text{ or there exist} \\ &\omega'_1, \omega'_2 \in \mathcal{B}_h(\omega) \text{ such that } h'a \prec h_1c_2 \lesssim \zeta(\omega'_1) \text{ and} \\ &h''b \prec h_2c_1 \lesssim \zeta(\omega'_2). \end{aligned} \tag{3}$$

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 \emptyset): two of them violate Condition 5 of Definition 6 (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 all causally restricted models.

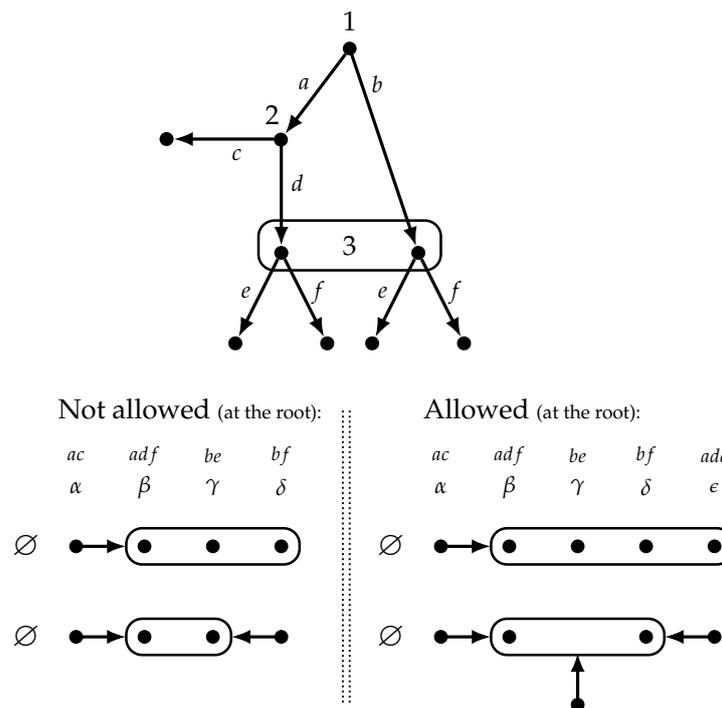


Figure 8. A game and four partial models of it (showing only the beliefs of Player 1 at history \emptyset), two of which violate Condition 5 of Definition 6 and the remaining two do not.

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 7. 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 8. 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 7), they predict the same choice(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 $h'_1 \prec h'a \prec \zeta(\omega_1)$ 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 $h'_2 \in [h_2], h'_2 \prec h''a \prec \zeta(\omega_2)$, and
2. if $\omega_2 \in \mathcal{B}_{h_2}(\omega)$ with $h'_2 \prec h''b \prec \zeta(\omega_2)$ 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 $h'_1 \in [h_1], h'_1 \prec h'b \prec \zeta(\omega_1)$.

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

$$\omega \in \mathbf{A} \text{ if and only if any two players active at } \omega \text{ are in agreement at } \omega. \tag{4}$$

We can now state our characterization result, according to which a self-confirming play is a Nash play if and only if the beliefs of any two players are in agreement about the hypothetical choice(s) of a third player at a future information set that they both consider. As in [1] we restrict attention to games that satisfy the property that each player moves at most once along any play. Equivalently, one could consider the agent form of the game, where the same player at different information sets is regarded as different players, but with the same payoff function.

Proposition 1. 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 1 is given in Appendix B.

Note that, in a perfect-information game, $\mathbf{T} \cap \mathbf{C} \subseteq \mathbf{A}$. Hence Proposition 1 in [1] is a corollary of the above Proposition.

4. Further Discussion and Conclusions

A reviewer suggested a discussion of the similarities and differences between the approach put forward in this paper and Steven Brams' Theory of Moves (TOM) [37] (see also the very recent [38]). TOM deals mostly with two-person strategic-form games in which each player has a strict ordinal ranking of the outcomes. TOM assumes that, instead of choosing strategies simultaneously and independently, players start from an outcome (that is, a strategy profile)—called the "initial state"—and from that outcome they consider the consequences of a series of moves and countermoves that lead from state to state. The sequence of moves and countermoves is strictly alternating and the process continues until the game terminates in a "final state" which is called the "final outcome" or simply "outcome" of the game. It is assumed that no payoffs accrue to players from being in a state unless it is the final state (which could be the initial state if the players choose not to move from it). Players make farsighted calculations of where play will terminate after a finite sequence of moves and countermoves. The result of such farsighted calculations

is called a Non-Myopic Equilibrium (NME). Thus, an NME can be understood as the backward-induction solution of a finite perfect-information “moves game” (with no ties).

There are two points in common between TOM and our approach. First of all, only ordinal preferences are considered in both approaches. Secondly, both approaches are based on a departure from standard solution concepts in game theory. However, the differences between the two approaches are substantial. Brams’ theory is not based at all on epistemic considerations: no beliefs are attributed to the players and the conceptual nature of a “state” is very different; in TOM a state is merely an outcome, while in our doxastic approach a state is *described* not only in terms of an outcome but also in terms of doxastic relations that describe what the active players believe about the possible outcomes when it is their turn to move. Furthermore, while TOM starts from a strategic-form game and builds on it a finite perfect-information game by specifying an initial outcome and the rules for moves and countermoves from it, we analyze a given extensive-form game (with possibly imperfect information) without modifying it in any way. Our approach falls within the “epistemic foundations approach” in which beliefs play an essential role. TOM, on the other hand, is entirely “belief-free”.

In Section 2, we raised the question “what constitutes a rational solution of an extensive-form game?” Most of the epistemic game theory literature has gone in the direction of imposing more and more subtle and complex conditions on counterfactual beliefs and choices of the players at unreached information sets. We suggested going in the opposite direction, by focusing only on the actions and beliefs of the active players. Within this framework, a natural notion of rational play is captured by the definition of self-confirming play, where each action taken is justified by the beliefs held at the time of choice and those beliefs turn out to be exactly correct, so that no player receives information that contradicts those beliefs and thus experiences no regret. This approach is flexible enough to allow one to explore the epistemic foundations of standard solution concepts such as Nash equilibrium and backward induction (within the behavioral approach described in this paper, the epistemic conditions needed to obtain a characterization of backward induction in perfect information games, or a generalized version of it for a class of games with imperfect-information, are investigated in [1] and [9], respectively).

The characterization of Nash play given in Proposition 1—unlike characterizations of Nash equilibrium provided for strategic-form games (for a discussion of the relevant literature the reader is referred to ([1], Section 6))—does not require players to believe in each other’s rationality. 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 Definition 4 we put forward the notion of self-confirming play, which is in the spirit of self-confirming equilibrium ([13]), but framed in behavioral terms and without making use of the notion of strategy. Whereas in perfect-information games the notion of self-confirming play is equivalent to the notion of Nash play, the equivalence does not extend to imperfect information games. Proposition 1 identified the additional restriction that is needed to characterize the set of Nash plays in games with imperfect information.

The main purpose of this paper was to show that one can go a long way in the analysis of rational play in extensive-form games without using the notion of strategy, that is, without the need to specify choices at all histories—even those that are not reached—and without the need to model players’ beliefs at unreached histories. We argued that the standard approach based on Nash equilibrium and its refinements is too ambitious in its goal to tackle the counterfactual behavior and beliefs of players at unreached histories and that there is no need to pursue this goal in order to have a theory of rational behavior in dynamic games.

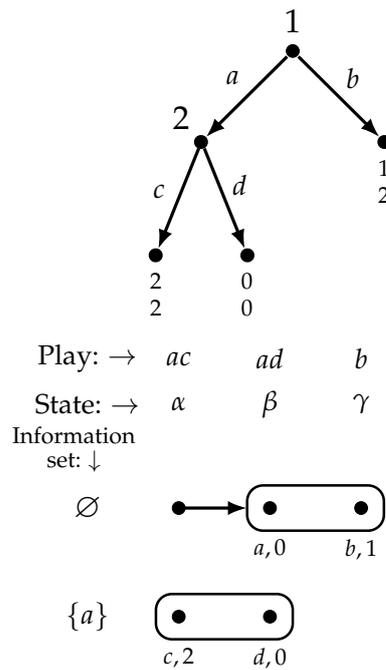


Figure 9. $\gamma \in R \cap T \cap C \cap A$ but at γ Player 1 does not believe that Player 2 is rational.

In what directions can the approach discussed in this paper be further developed? One natural extension is to move from ordinal preferences to von Neumann-Morgenstern preferences and from qualitative beliefs to probabilistic beliefs; one would then, correspondingly, move from the very weak definition of rationality given in Definition 2 to the stronger definition of rationality as expected utility maximization. Another possible line of inquiry would be to identify the circumstances (if any) that would make the structures used in this paper inadequate and would require a full analysis in terms of counterfactuals (which in turn would require extending those structures by adding similarity relations among states, as explained in Section 3).

Funding: This research received no external funding.

Data Availability Statement: Not applicable.

Acknowledgments: I am grateful to three anonymous reviewers and to participants in the Workshop on Epistemic Game Theory (EPICENTER, Maastricht University, July 2022) and the LOFT conference (Groningen University, July 2022) for useful comments.

Conflicts of Interest: The author declares no conflict of interest.

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

A *finite extensive form without chance moves* is given by the following elements, where all the sets are finite:

1. A set of players denoted by N .
2. A set of actions, denoted by A .

3. A set of histories, denoted by $H \subseteq A^*$, which satisfies the property that, 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 , belongs to 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*. Z denotes the set of terminal histories and $D = H \setminus Z$ the set of *decision* histories.
4. To every decision history is assigned a player, by means of a function $\iota : D \rightarrow N$. Thus, $\iota(h) \in N$ is the player who moves, or *is active*, at $h \in D$. For notational simplicity we assume that there is exactly one player who is active active at any decision history; thus, a simultaneous move by, say, Players 1 and 2 is represented in the traditional way by having Player 1 move first followed by Player 2, who is not informed of Player 1's move. Let $D_i = \{h \in D : i = \iota(h)\}$ denote the set of histories at which player i is active. For every $h \in D$, $A(h)$ denotes 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 $i \in N$, we postulate an equivalence relation \approx_i on D_i : $h \approx_i h'$ if and only if, 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 actions available at an information set are not allowed to differ across histories in that information set, that is, if $h \approx_i h'$ then $A(h') = A(h)$. 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_1a 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' .

When every information set consists of a single history, the game is said to have *perfect information*, otherwise it is said to have *imperfect information*.

In order to lighten the notation, 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 $a_1a_2a_3a_4$.

An *extensive game with ordinal payoffs* is obtained from a given extensive form, 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 z to be at least as good as z' . It is often convenient to replace the relation R_i with a real-valued *utility* (or *payoff*) function $u_i : Z \rightarrow \mathbb{R}$ satisfying the property that $u_i(z) \geq u_i(z')$ if and only if $z R_iz'$.

Appendix B. Proof of Proposition 1

Given a finite extensive-form game and a pure-strategy profile s , define the function $f_s : H \rightarrow Z$ as follows: if $z \in Z$ (that is, if z is a terminal history) then $f_s(z) = z$ and if $h \in D$ (that is, if h is a decision history) 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 $h' \in [h]$, $h' \not\prec z_s^*$. 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 A1. 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 $g_s : I \rightarrow D$ 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 A2. 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.

- $\Omega = Z$.
- $\zeta : Z \rightarrow Z$ 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 $h \not\prec z$, then $\mathcal{B}_h(z) = \emptyset$.

2. If $h \prec z_s^*$ then $\mathcal{B}_h(z_s^*) = \{z' \in Z : z' = f_s(ha) \text{ for some } a \in A(h)\}$. [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 $h \not\prec z_s^*$, but $[h]$ is not avoided by z_s^* , then, for all $z \in Z$ such that $h \prec z$, $\mathcal{B}_h(z) = \{z' \in Z : z' = f_s(h_s^*([h])a) \text{ for some } a \in A(h)\}$. [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 $h' \prec z$, $\mathcal{B}_{h'}(z) = \{z' \in Z : z' = f_s(\hat{h}a) \text{ for some } a \in A(h)\}$. [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 A1. Note that the model generated by a pure-strategy profile s and a selection function g_s is a causally restricted model (Definition 6).

Remark A2. 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 A2). Then the no-uncertainty conditions 1 and 2 of Definition 3 and the agreement condition (4) are satisfied at every state, that is, $\mathbf{C} = \mathbf{A} = Z$. Furthermore, by Point 1 in Definition A2, $z_s^* \in \mathcal{B}_h(z_s^*)$ for all h such that $h \prec z_s^*$; that is, $z_s^* \in \mathbf{T}$.

We can now prove Proposition 1.

Proof. (A) [Note that, for this part of the proof, the restriction that no player moves more than once along any play of the game is not needed.] 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 A1) and consider the model generated by s and g_s (Definition A2). By Remark A2, $z_s^* \in \mathbf{C} \cap \mathbf{T} \cap \mathbf{A}$ (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 z_s^* (that is, $h \prec z_s^*$) and let a be the action at h such that $ha \succ z_s^*$, that is, $s(h) = a$; then $f_s(ha) = f_s(\emptyset) = z_s^*$. Suppose that player $\iota(h)$ is not rational at h . Then there must be a $b \in A(h) \setminus \{a\}$ that guarantees a higher utility to player $\iota(h)$: if $z' \in \mathcal{B}_h(z_s^*)$ is such that $hb \succ z'$, then $u_{\iota(h)}(z') > u_{\iota(h)}(z_s^*)$. By Definition A2, $z' = f_s(hb)$ so that $u_{\iota(h)}(f_s(hb)) > u_{\iota(h)}(f_s(ha))$; hence, by unilaterally changing her strategy at h from a to b (while leaving the rest of her strategy unchanged), player $\iota(h)$ can increase her payoff, contradicting the assumption that s is a Nash equilibrium.

(B) Fix a finite extensive-form game G where no player moves more than once along any play and consider an arbitrary 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 want to show that we can construct a pure-strategy Nash equilibrium s of G such that $f_s(\emptyset) = \zeta(\alpha)$.

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

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 is not a prefix of $\zeta(\alpha)$ (that is, $b \neq s(h)$ where $s(h)$ was defined in Step 1). By Definition of model (Definition 1) there exists an $\hat{\omega} \in \mathcal{B}_h(\alpha)$ such that $\hat{h}b \succ \zeta(\hat{\omega})$ for some $\hat{h} \in [h]$. Since $\alpha \in \mathbf{C}$, by Point 1 of Definition 3 for every $\omega' \in \mathcal{B}_h(\alpha)$ and for every $h' \in [h]$, if $h'b \succ \zeta(\omega')$ then $h' = \hat{h}$. Since $\alpha \in \mathbf{C}$, by Point 2 of Definition 3 for any other $\omega \in \mathcal{B}_h(\alpha)$ such that $\hat{h}b \not\succ \zeta(\omega)$, $\zeta(\omega) = \zeta(\hat{\omega})$. Define, for every h' such that $\hat{h}b \not\succ \zeta(\hat{\omega})$, $s(h') = c$ where $c \in A(h')$ is the action at h' such that $h'c \succ \zeta(\hat{\omega})$. 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.

Steps 1 and 2 define the choices prescribed by s along the play $\zeta(\alpha)$ as well as 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 $h \preceq \zeta(\alpha)$ (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 such that $h \prec \zeta(\alpha)$ (that is, h reached at state α) and, by switching her choice at h from $s(h)$ to a different choice, player $i(h)$ can increase her payoff (by hypothesis there are no successors of h that belong to player $i(h)$). Let $s(h) = a$ (that is, $ha \preceq \zeta(\alpha)$) and let b be the choice at h that yields a higher payoff to player $i(h)$; that is,

$$u_{i(h)}(f_s(hb)) > u_{i(h)}(\zeta(\alpha)). \quad (\text{A1})$$

By Item 4 of Definition 1 there exists a $\beta \in \mathcal{B}_h(\alpha)$ such that $hb \preceq \zeta(\beta)$. Since $\alpha \in \mathbf{C}$, for every $\omega \in \mathcal{B}_h(\alpha)$ such that $hb \preceq \zeta(\omega)$, $\zeta(\omega) = \zeta(\beta)$. By Step 2 above,

$$\zeta(\beta) = f_s(hb). \quad (\text{A2})$$

Hence, by (A2), at decision history h and state α , player $i(h)$ believes that if she plays b her payoff will be $u_{i(h)}(f_s(hb))$. 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 $ha \preceq \zeta(\omega)$, $\zeta(\omega) = \zeta(\alpha)$. Hence, at state α and history h , player $i(h)$ believes that if she plays a her payoff will be $u_{i(h)}(\zeta(\alpha))$. It follows from this and (A1) that at α and h player $i(h)$ believes that action b is better than action a (Definition 2), which implies that at α player $i(h)$ is not rational, contradicting the assumption that at $\alpha \in \mathbf{R}$. \square

References

- Bonanno, G. Behavior and deliberation in perfect-information games: Nash equilibrium and backward induction. *Int. J. Game Theory* **2018**, *47*, 1001–1032. [CrossRef]
- Kreps, D.; Wilson, R. Sequential equilibrium. *Econometrica* **1982**, *50*, 863–894. [CrossRef]
- Lewis, D. *Counterfactuals*; Harvard University Press: Cambridge, MA, USA, 1973.
- Stalnaker, R. A theory of conditionals. In *Studies in Logical Theory*; Rescher, N., Ed.; Blackwell: Oxford, UK, 1968; pp. 98–112.
- Stalnaker, R.; Thomason, R. A semantical analysis of conditional logic. *Theoria* **1970**, *36*, 246–281.
- Perea, A. Backward induction *Versus* Forw. Induction Reason. *Games* **2010**, *1*, 168–188. [CrossRef]
- Perea, A. Belief in the opponents' future rationality. *Games Econ. Behav.* **2014**, *83*, 231–254. [CrossRef]
- Baltag, A.; Smets, S.; Zvesper, J. Keep 'hoping' for rationality: A solution to the backward induction paradox. *Synthese* **2009**, *169*, 301–333.
- Bonanno, G. A doxastic behavioral characterization of generalized backward induction. *Games Econ. Behav.* **2014**, *88*, 221–241.
- Battigalli, P.; Siniscalchi, M. Strong belief and forward induction reasoning. *J. Econ. Theory* **2002**, *106*, 356–391.
- Pearce, D. Rationalizable strategic behavior and the problem of perfection. *Econometrica* **1984**, *52*, 1029–1050.
- Meier, M.; Perea, A. Forward Induction in a Backward Inductive Manner. Technical Report, EPICENTER, Maastricht University. 2022. Available online: <https://www.epicenter.name/Perea/Papers/BI-FI-procedure.pdf> (accessed on 23 September 2022).
- Fudenberg, D.; Levine, D. Self-confirming equilibrium. *Econometrica* **1993**, *61*, 523–545.
- Battigalli, P.; Guaitoli, D. Conjectural equilibria and rationalizability in a game with incomplete information. In *Decisions, Games and Markets*; Battigalli, P., Montesano, A., Panunzi, F., Eds.; Kluwer Academic Publishers: Dordrecht, The Netherlands, 1997; pp. 97–124.
- Greenberg, J. The right to remain silent. *Theory Decis.* **2000**, *48*, 193–204.
- Greenberg, J.; Gupta, S.; Luo, X. Mutually acceptable courses of action. *Econ. Theory* **2009**, *40*, 91–112.
- Dekel, E.; Fudenberg, D.; Levine, D. Payoff information and self-confirming equilibrium. *J. Econ. Theory* **1999**, *89*, 165–185. [CrossRef]
- Dekel, E.; Fudenberg, D.; Levine, D. Subjective uncertainty over behavior strategies: A correction. *J. Econ. Theory* **2002**, *104*, 473–478. [CrossRef]
- Fudenberg, D.; Kamada, Y. Rationalizable partition-confirmed equilibrium. *Theor. Econ.* **2015**, *10*, 775–806.
- Battigalli, P. *Comportamento Razionale ed Equilibrio nei Giochi e nelle Situazioni Strategiche*. Unpublished Dissertation, Bocconi University, Milano, Italy, 1987.

21. Gilli, M. Metodo Bayesiano e Aspettative nella Teoria dei Giochi e nella Teoria Economica. Unpublished Dissertation, Bocconi University, Milano, Italy, 1987.
22. Halpern, J. Substantive rationality and backward induction. *Games Econ. Behav.* **2001**, *37*, 425–435.
23. Samet, D. Hypothetical knowledge and games with perfect information. *Games Econ. Behav.* **1996**, *17*, 230–251. [[CrossRef](#)]
24. Battigalli, P.; Di-Tillio, A.; Samet, D. Strategies and interactive beliefs in dynamic games. In *Advances in Economics and Econometrics. Theory and Applications: Tenth World Congress, Volume 1*; Acemoglu, D., Arellano, M., Dekel, E., Eds.; Cambridge University Press: Cambridge, UK, 2013; pp. 391–422.
25. Bonanno, G. A dynamic epistemic characterization of backward induction without counterfactuals. *Games Econ. Behav.* **2013**, *78*, 31–43.
26. Bonanno, G. The material conditional is sufficient to model deliberation. *Erkenntnis*, February 2021. ISSN 1572-8420. Available online: <https://link.springer.com/article/10.1007/s10670-020-00357-7> (accessed on 23 September 2022).
27. Stalnaker, R. Knowledge, belief and counterfactual reasoning in games. *Econ. Philos.* **1996**, *12*, 133–163.
28. Aumann, R. Backward induction and common knowledge of rationality. *Games Econ. Behav.* **1995**, *8*, 6–19. [[CrossRef](#)]
29. Aumann, R. On the centipede game. *Games Econ. Behav.* **1998**, *23*, 97–105. [[CrossRef](#)]
30. Stalnaker, R. Belief revision in games: Forward and backward induction. *Math. Soc. Sci.* **1998**, *36*, 31–56. [[CrossRef](#)]
31. Bicchieri, C.; Green, M. Symmetry arguments for cooperation in the Prisoner’s Dilemma. In *The Logic of Strategy*; Bicchieri, C., Jeffrey, R., Skyrms, B., Eds.; Oxford University Press: Oxford, UK, 1999; pp. 175–195.
32. Gauthier, D. *Morals by Agreement*; Oxford University Press: Oxford, UK, 1986.
33. Nozick, R. Newcomb’s problem and two principles of choice. In *Essays in Honor of Carl G. Hempel: A Tribute on the Occasion of His Sixty-Fifth Birthday*; Rescher, N., Ed.; Springer: Dordrecht, The Netherlands, 1969; pp. 114–146. [[CrossRef](#)]
34. Spohn, W. Dependency equilibria and the causal structure of decision and game situations. *Homo Oeconomicus* **2003**, *20*, 195–255. [[CrossRef](#)]
35. Spohn, W. Dependency equilibria. *Philos. Sci.* **2007**, *74*, 775–789. [[CrossRef](#)]
36. Spohn, W. From Nash to dependency equilibria. In *Logic and the Foundations of Game and Decision Theory—LOFT 8*; Bonanno, G., Löwe, B., van der Hoek, W., Eds.; Springer: Berlin/Heidelberg, Germany, 2010; pp. 135–150.
37. Brams, S.J. *Theory of Moves*; Cambridge University Press: Cambridge, UK, 1993.
38. Brams, S.; Ismail, M. Every normal-form game has a Pareto-optimal nonmyopic equilibrium. *Theory Decis.* **2022**, *92*, 349–362. [[CrossRef](#)]