

NOTICE WARNING CONCERNING COPYRIGHT RESTRICTIONS:

The copyright law of the United States (title 17, U.S. Code) governs the making of photocopies or other reproductions of copyrighted material. Any copying of this document without permission of its author may be prohibited by law.

Forward Induction

by

Gian Aldo Antonelli and Cristina Bicchieri

June 1994

Report CMU-PHIL-58



**Philosophy
Methodology
Logic**

Pittsburgh, Pennsylvania 15213-3890

**University Libraries
Carnegie Mellon University
Pittsburgh PA 15213-3890**

Forward Induction*

Gian Aldo Antonelli
Yale University

Cristina Bicchieri
Carnegie Mellon University

June 1994

Abstract

In this paper we isolate a particular refinement of the notion of Nash equilibrium that is characterized by two properties: (i) it provides a unified framework for both backwards and forward induction; and (ii) it is mechanically computable. We provide an effective procedure that allows players, given the extensive-form representation of a game, to compute a set of “reasonable paths” through the tree. The set of reasonable paths corresponds to the set of strategies that survive iterated elimination of (weakly) dominated strategies in the strategic form. We prove that whenever our procedure identifies a unique path, that path corresponds to a Nash equilibrium. Moreover, our procedure rules out all Nash equilibria that contain (weakly) dominated strategies. Further, our notion of “reasonable” paths leads to the backwards induction solution in the case of games of perfect information, and to forward induction in the case of games of imperfect information. We model the players’ reasoning process by giving a theory (with which each player is supposed to be endowed), from which statements characterizing the players’ behavior are deducible. Such a theory is not yet *complete*, in that it cannot handle true (irrational) deviations. We point at directions for future work by showing how such a theory can be made complete provided we reinterpret some of its axioms as defeasible inference rules.

1 Introduction

Two problems have been widely discussed in recent debates on the foundations of game theory, namely: (i) the problem of how to handle deviations from equilibrium play, and (ii) the problem of how to model counterfactual reasoning patterns that may be necessary to justify solutions such as backwards induction (Aumann [1], Stalnaker[13]). There are obvious connections between these two problems, although they have usually been treated separately.

The first problem appears in the context of games that have multiple Nash equilibria, some of which might be implausible in that they involve *risky* (i.e. weakly dominated) strategies and the implausible beliefs that those strategies will be played. Various refinements of Nash equilibrium have been proposed to take care of implausible equilibria, as well as to attain predictability in the face of multiplicity. In games of perfect information this is accomplished by applying backwards induction. In games of imperfect information instead, one has to appeal to different types of refinements, each of which corresponds to a different way of checking the stability of a Nash equilibrium against deviations from equilibrium play. The stability of an equilibrium, however, is a function of how a deviation is being interpreted. Counterfactuals play a role in this context since, from the viewpoint of a particular equilibrium, an off-equilibrium move is a contrary-to-fact event. When considering the possibility of such an event, a player has to undergo a belief-revision process, retracting from

*An earlier version of this paper appeared in [6, pp. 24–43].

his original set of beliefs all those beliefs that contradict the statement that an off-equilibrium event has taken place. The interpretation of out-of-equilibrium play will thus depend on the model of belief revision adopted (and on the interpretation of the counterfactual), and so will the resulting refinement (Bicchieri [2, 3, 4]).

As a refinement for games of perfect information, backwards induction embodies the principle that a rational player will only play undominated strategies. In this context, too, the issue of counterfactuals arises when a player contemplates a deviation from equilibrium play. Though thinking about possible deviations is not required if one simply wants to compute a solution, an analysis of deviations becomes crucial when justifying a given solution. If the theory of the game is expressed in a monotonic logic, then the union of that theory with a statement to the effect that a move outside the solution has taken place will be inconsistent. To regain consistency, one may want to provide the players with a belief-revision model, or include counterfactuals into the players' own theory of the game (Bicchieri & Antonelli [5]). Augmenting the theory of the game with an account of counterfactuals can solve the problem, at a price. Different kinds of game may need different accounts of counterfactuals to accommodate deviations, whereas one would like to have a unique general account of counterfactuals that applies to all types of games.

It is often argued that a complete theory of the game is a theory that explains the unexpected, that is, it is a theory that explains all sorts of moves, including irrational ones, on the part of the players. The task of constructing such theories, however, has proved quite formidable. For example, a theory that interprets deviations as mistakes takes into account all sorts of deviations, but, as we shall show in Section 2, a theory of mistakes may be incompatible with rationality being common knowledge. On the other hand, a theory that only considers rational deviations, interpreting them as signals, is consistent with rationality being common knowledge, but it is silent as far as irrational deviations are concerned. In the present paper, we shall argue that the treatment of deviations, be they studied in finite extensive form games of perfect or imperfect information, does not require a full account of (intra-theoretical) counterfactual reasoning, notoriously one of the thorniest issues in philosophical logic and knowledge representation.

In our model, the players reason to a solution on the basis of the theory of the game they are endowed with. Players are like automatic theorem provers provided with a decision procedure that isolates, whenever possible, a unique solution that satisfies the rationality conditions embedded in the theory's axioms. A surprising result of such a decision procedure is that it leads quite naturally to the backwards induction equilibrium in finite games of perfect information, and to the forward induction refinement in finite games of imperfect information.

The theory of the game T is formalized in classical first-order logic and is a revised version of Bicchieri & Antonelli [5]. Our theory, which is interpretable in Primitive Recursive Arithmetic, comprises general axioms describing the game (represented by a finite tree) and the payoff structure for each player. We supply a function x^* giving, for each information set, the set of undominated paths starting at that set. Finally, we give "behavioral" axioms describing how players' actions are determined by the set of undominated paths and hence, indirectly, by their expected payoff at each information set. This theory allows players to infer the sequence of moves comprising what we call a Reasonable Path (or paths), i.e. a path that satisfies the rationality conditions that are embedded in our theory.

The theory we propose here differs from Bicchieri & Antonelli [5] in several important respects. It is far more general, as it also includes games of imperfect information, and it is not assumed that players have "local knowledge" at an information set; in fact, it is assumed that the theory T is group knowledge among the players¹. Since we model a decision procedure that leads to the

¹By group knowledge of p we mean that every member of the group knows p .

Figure 1: A two-person game of imperfect information.

isolation of a subset of undominated paths, no belief revision is needed to identify the subset of Nash equilibria (if such a subset exists) that contains only undominated strategies. Belief revision is only necessary when a player deviates to a dominated path. In our view, it makes sense to call "deviation" only an action that is unexpected in that it is obviously contrary to the best interest of the deviant player. Such action should then be explained.

The theory T we assign to the players only provides a background description of the game. Considering a deviation means augmenting T by a "history," i.e. by axioms A_1, \dots, A_k specifying that certain moves have been made. Since those moves lie outside the undominated path(s), $T + A_1 + \dots + A_k$ is inconsistent. This implies that any model of belief revision based on T and meant to accommodate deviations cannot but generate a modification of T .

By exhibiting theory T we accomplish a twofold task: First, we show how a first-order theory of the game is perfectly adequate to infer the backwards induction equilibrium in games of perfect information and, in games of imperfect information, it gives us a refinement that agrees with forward induction. These results are obtained without an account of counterfactuals. Second, we set the stage for a default version of the theory obtained by the original theory T by weakening the axioms, i.e. by reinterpreting certain material implications as defeasible default rules. A default theory of the game is a complete theory in the sense that it can be augmented with information to the effect that a *true* deviation (i.e. an irrational move) from an equilibrium path has occurred without becoming inconsistent.

2 An Example

In the refinements literature, anticipated actions off the equilibrium path play a crucial role in sustaining an equilibrium. Given a Nash equilibrium, the players are supposed to ask what would happen if one of them were to deviate from the equilibrium path, and an equilibrium is considered plausible (or stable) only in case the players would have no incentive to play another strategy in face of a deviation. From the viewpoint of a given Nash equilibrium, asking what to do when a deviation occurs is tantamount to asking a counterfactual question, since an off-equilibrium move is by definition a contrary-to-fact event. Though most game theorists now recognize that the treatment of deviations involves counterfactual reasoning and a change of beliefs on the part of the players, there are very few syntactical or semantical models of belief revision in the literature. Furthermore, little attention has been paid to whether the beliefs attributed to the players are reasonable, in the sense of being consistent with their information about the game. What follows is an example of how the problem of multiple equilibria is usually addressed, where the kinds of refinement proposed depend upon the beliefs attributed to the players.

The game in Figure 1 is a two-person game of imperfect information. As usual, the players are assumed to have common knowledge of rationality (i.e. that they are expected utility maximizers) and of the structure of the game. Player 1 has three choices: either he chooses c , which ends the game, or he may choose b or a , in which case it is player 2's turn to move. If player 2 is called upon to move, however, she will not know player 1's preceding move, so she cannot tell whether she is at node y or z . The game has two Nash equilibria in pure strategies, (c, L) and $(a, J2)$, and one would like to find some means to predict which equilibrium will be chosen by the players. Yet in such a simple game the most common refinement concepts, such as perfection, properness, or sequential equilibrium do not succeed in selecting a unique equilibrium. Let us see why.

The equilibria (c, I) and (a, R) are both perfect (Selten [12]). In particular, (c, Z) is perfect if player 2 believes that 1 will make mistake b with a greater probability than mistake a , but whereas both probabilities are very small, the probability of playing c is close to 1. If this is player 2's belief, then she should play L with probability close to 1. In this case, player 1 should play c . But why should 2 believe that mistake b is more likely than mistake a ? Since strategy c strictly dominates b , there is no reason to expect mistake b to occur more frequently than mistake a . The beliefs that support (c, L) are thus unreasonable.² The problem is that out-of-equilibrium beliefs are unrestricted: A player is supposed to ask whether it is reasonable to believe that the opponent will play his part in a given Nash equilibrium, but not whether the beliefs supporting the opponent's choice are rational. A rational belief, in this case, means a belief consistent with rationality being common knowledge. In our example, player 1 attributes to player 2 a belief that justifies her choice of strategy i , but it is not obvious that 2's belief about the greater likelihood of mistake b is defensible.

It could be argued that one way to restrict out-of-equilibrium beliefs is to restrict a player's conjectures about the opponent's behavior to those that are rationally justified. A rational player, for example, should be expected to avoid costly mistakes (Myerson [9]). Proper equilibria need only be robust with respect to plausible deviations, i.e., deviations that do not involve costly mistakes. However, one mistake may be more costly than another only insofar as the player who could make the mistake has definite beliefs about the opponent's reaction. In our example, both (c, ξ) and (a, R) are proper equilibria. If a deviation from (c, L) were to occur, player 2 would keep playing L only if she were to assign a higher probability to mistake b than to mistake a . And if player 1 were to expect 2 to play i , mistake b would indeed be less costly than a . In this case, L would be a best reply for player 2. Thus mistake b is less costly if 1 expects 2 to play Z , and 2 will play L only if she believes that 1 expects her to play L in response to a deviation. But, again, why should player 2 be expected to play L in the first place? Since b is strictly dominated by c , it is very unlikely that b occurs. The only plausible deviation is thus a , but then player 2 should be expected to play R .

The same problem arises with the sequential equilibrium notion (Kreps & Wilson [8]), which explicitly specifies beliefs at information sets lying off the equilibrium path. In our example, both (c, Z) and (a, R) are sequential equilibria, since an equilibrium strategy has to be optimal with respect to some beliefs, but not necessarily plausible beliefs. In particular, if player 1 chooses c , then any probability assessment by player 2 is reasonable, and it is entirely possible that player 2 assesses a higher probability to strategy b than to a .

The equilibrium (c, L) is intuitively unreasonable precisely because the beliefs that support it are unreasonable. By reasonable beliefs we mean beliefs that are consistent with a player's background beliefs and knowledge. Since in our example players have common knowledge of rationality, a player should never be expected to choose a dominated strategy. This must be true of weakly dominated strategies, too. The rationale for this requirement is simple: Since off-equilibrium choices are relevant only when they affect the choices along the equilibrium path, it is reasonable to ask that an off-equilibrium choice that is weakly dominated should be ruled out, since it is as good as some other strategy if the opponent sticks to the equilibrium, but it does worse when a deviation occurs. In our example, rationality is common knowledge and strategy b is strictly dominated, therefore it must be common knowledge that b is never going to be played. Knowing that player 2 will always respond to a deviation with iZ , player 1 will have an incentive to choose a . This kind of reasoning rules out (c, L) as implausible.

Considering only undominated choices means that off-equilibrium beliefs should satisfy the

²Even Selten [12, p. 35] admits that game theory is concerned with absolutely rational agents and that "there cannot be any mistakes if the players are absolutely rational."

Figure 2: The strategic form of the game of Figure 1.

following condition:

- (R) When considering an off-equilibrium move, a player should not hold beliefs that are inconsistent with common knowledge of rationality.

All that condition (R) tells us is that whenever a player has a weakly dominated strategy he should not be expected to use it, and that no one should choose a strategy that is a best reply to an opponent's weakly dominated strategy. Common knowledge of rationality thus implies common knowledge that weakly dominated strategies will not be played. Note that condition (R) entails iterated elimination of dominated strategies in the strategic form. Consider for example the strategic form of the game in Figure 1, which is given in Figure 2. In this game b is eliminated since it is strictly dominated by c . Since b is eliminated, R weakly dominates L , which is in turn eliminated. Finally, a dominates c for player 1, hence (a, R) is the only equilibrium that survives iterated elimination of dominated strategies.

Is there a correspondence between the iterated procedure we have just described and our informal argument in favor of (a, R) in the extensive form? If the game in Figure 1 were one of perfect information, backwards induction would give us a decision procedure that matches iterated elimination in the strategic form. Starting from terminal nodes, players eliminate weakly dominated strategies bottom up; in the absence of ties, this method determines a single outcome. In our example, it would be the equilibrium (a, R) . Note that backwards induction requires rational behavior even in those parts of the tree that may not be reached if an equilibrium is played. As a result, backwards induction leads to eliminate all but the equilibrium points that are in equilibrium in each of the subgames and in the entire game. More generally, we may state the following backwards induction condition:

- (BI) A strategy is optimal only if that strategy is optimal when the play begins at any information set that is not the root of the game tree.

In games of perfect information, (R) and (BI) guarantee that unreasonable equilibria are ruled out. Together, they imply that a plausible Nash equilibrium must be consistent with deductions based on the opponent's rational behavior in the future. Future behavior, however, may involve off-equilibrium play, and in this case condition (R) tells us that the only deviations that matter are undominated choices, i.e., choices that can be interpreted as intentional moves of rational players.

The game in Figure 1, however, is one of imperfect information; here the backwards induction algorithm fails because it presumes that an optimal choice exists at every information set, given a specification of play at the successors. At 2's information set, however, there is no unique rational action: at node y she should play L , and at node y' she should play R . Even if backwards induction is not defined, conditions (R) and (BI) may still apply to games of imperfect information that have proper subgames. For each such subgame, one may ask whether an equilibrium for the whole game induces an equilibrium in the subgame (Selten [11]). Yet the game in Figure 1 has no proper subgames, so in this case (BI) does not apply. Condition (R) still applies, though, by constraining the possible interpretations of deviations.

In Figure 1, for example, if player 2 gets to play then player 1 must have foregone the sure payoff of 3 in favor of playing a . The only equilibrium that yields a payoff greater than 3 to player 1 is (a, R) , hence 2 should deduce from the fact that her information set is reached that 1 has

Figure 3: A more complex two-person game.

Figure 4: The strategic form of the game in Figure 3.

chosen strategy a . If so, 2's best reply is R and player 1, anticipating player 2's reasoning, will conclude that it is optimal for him to play a . What we have just described is a forward induction argument, that is, an argument based on inferences about the opponent's rational behavior in the past. In our example there is no past to speak of, but rather the knowledge that player 1, facing the choice of getting a payoff of 3 for sure or playing a simultaneous game with player 2, has chosen the second option. A forward induction argument thus interprets deviations from a given equilibrium as signals, intentional choices of a rational player (Kohlberg & Mertens [7]). For this interpretation of deviations to be consistent with rationality, however, there must exist at least a strategy that yields the deviating player a payoff greater or equal to that obtained by playing the equilibrium strategy. This consideration leads to the following iterated dominance requirement:

(ID) A plausible equilibrium of a game G must remain a plausible equilibrium of any game G' obtained from G by deleting (weakly) dominated strategies.

Condition (ID) implies the iterated use of condition (R) in games that have subgames. Taken together, conditions (ID), (R) and (BI) underlie the forward induction argument.

Consider the following game: In Figure 3 each player has the choice of playing down, which ends the game, or playing across. At node w , if player 1 chooses to play across he plays a simultaneous battle of the sexes with player 2. This game has two equilibria in pure strategies: (A_1A_3T, D_2) and (A_1Z_3, A_2R) . Note that in the strategic form of Figure 3 the equilibrium (A_1Z_3, A_2R) does not survive iterated elimination of (weakly) dominated strategies.

Like other refinements, forward induction is used to check the equilibria of the game against possible deviations. The difference with other refinements lies in the criteria used to assess deviations. Suppose the players agree to play (A_1Z_3, A_2R) but, unexpectedly, player 1 deviates at node w by playing A_3 . Since A_3B is dominated by $\$3$, condition (R) rules out A_3B . Player 2 will then know that, if her information set is reached, A_3T has been played, therefore she will respond with L . Foreseeing this reasoning of player 2, player 1 should play A_3T . In the subgame G' starting at node w , the equilibrium profile $\{D_2, B^R\}$, though subgame perfect, is ruled out by a forward induction argument. The equilibrium $\{A_1, D^A_2R\}$ thus violates all three conditions (R), (BI) and (ID).

Consider now equilibrium (A_1A_3T, Z_2) . What happens if a deviation occurs at node y ? Condition (R) suggests that, by deviating, player 2 expects a higher payoff than what she gets by playing Z_2 . So it must be the case that, by deviating from the equilibrium path, player 2 is signaling that she will play R in the battle of the sexes. In which case it would be better for player 1 to play D_3 . However, this reasoning is fallacious, since condition (R) must be applied iteratively. If rationality is common knowledge, it must also be common knowledge that, at node w , player 1 will not play D_3 but A_3T . Hence 2's best reply is L . Since it must be common knowledge that at node w player 1 will play A_3T , it follows that in the subgame G' starting at node y player 2 will choose Z_2 . The equilibrium (A_1A_3T, Z_2) survives iterated application of condition (R), and in addition it also satisfies (BI) and (ID).

Given the above argument it follows that, if rationality is common knowledge, deviation A_2 should never be observed. Thus if the forward induction refinement has the advantage of making

off-equilibrium beliefs consistent with common knowledge of rationality, its drawback is that it does not provide a complete theory of the game: Unexpected deviations cannot happen. One way to address this issue is to *complete* the theory with a model of belief revision, but in this paper we wish to take a different path.

Instead of computing the Nash equilibria for the game and then test them for stability against deviations, we model how the players themselves may reason to a solution. The nature of such solutions will obviously be a function of the way in which we characterize the players and their information. In the next section, we provide the players with a theory of the game T (a set of axioms) that embodies a simple rationality condition. We show that T leads to iterative elimination of (weakly) dominated paths, that is, T generates an automatic decision procedure for the extensive form corresponding to iterated elimination of (weakly) dominated strategies in the strategic form. Through iterated elimination of (weakly) dominated paths we obtain a set (hopefully, a singleton) of undominated branches. We prove that whenever there exists a unique such branch, that branch corresponds to a pure strategy Nash equilibrium for the game. Moreover, all solutions thus obtained satisfy the criteria (R), (BI) and (ID). Note that the converse may not be true: if the game has a unique pure strategy Nash equilibrium that contains weakly dominated strategies, our procedure rules it out. In the last section, we show that for T to be a complete theory it is sufficient to interpret some of its axioms as default rules.

3 The Theory

In this section we are going to define a mechanical procedure that computes the set of undominated paths through a tree representing a game of imperfect information in extensive form. The procedure is defined by recursion on the height of the tree: for each node, the procedure computes the set of paths that are undominated in the subtree starting at that node.

Before giving the formal definitions, we are going to explain the ideas behind the procedure. First of all, we distinguish between (weakly) undominated *nodes* and (weakly) undominated *moves*. A node is weakly dominated within the context of a theory of the game if the theory predicts that that node will never be reached. For instance, in a game of perfect information, given certain assumptions of rationality, a node x (weakly) dominates a node y for player i (relative to a set P of paths) if every path $s \in P$ through x gives i a payoff at least as good as that of any path $s' \in P$ through y , and there is at least a path $s \in P$ through x that gives i a payoff that is strictly better than any payoff given by a path $s' \in P$ through y . All these notions will be given precise definitions below.

On the other hand, a move is weakly dominant at an information set if playing that move at any node of the information set guarantees a payoff that is not worse than any payoffs obtained by any other move and there is at least a payoff reachable playing that move that is strictly better than any payoff reachable playing any other move. Again, these ideas will be precisely defined below. Notice however that “reachable” is a notion that is relative to a given set P of paths: reachable from a node means reachable by a path $s \in P$ passing through the node.

The computing procedure comprises a *basic module* that is applied recursively at each information set. Our basic module is essentially a two-pass subroutine which is applied to a given information set I and returns a set of paths beginning at nodes in the information set. Call a node final if moving at that node ends the game. We can describe the behavior of the subroutine as follows. As a preliminary step, the subroutine computes a set P of paths – the paths under consideration: If the information set I given as input contains only final nodes then P contains all paths originating at those nodes, whereas if the nodes are not final, P is obtained by means of

Figure 5: An unbalanced tree.

recursive calls of the subroutine on the (information sets of) children of nodes in I .

In the first pass we eliminate from the set P of paths under consideration any paths beginning at nodes that are (weakly) dominated (relative to P). The first pass allows us to eliminate certain paths based on the position of their first node in the overall tree representing the game; such a first pass corresponds to forward induction. Let P' be the subset of P comprising those paths that survive the first pass. Then the second pass selects, from the paths in P' , those that corresponds to undominated moves.

When applying the procedure to the entire tree, we obtain a set of paths through the tree that are undominated and have the property that their restriction to any subtree still gives a set of undominated paths in the subgame represented by the subtree. Observe also that our procedure does not distinguish between games of perfect and imperfect information: the former are regarded as a limiting case of the latter, one in which all information sets are singletons.

We can now proceed with the formal details. A generic finite extensive form game of imperfect information G is represented by a finite tree, having an arbitrary branching factor, equipped with two functions p and I . Function $p : G \rightarrow \{1, \dots, k\}$ assigns a player i (for $0 < i \leq k$) to each node, while $I : G \rightarrow \mathcal{P}(G)$ assigns to each node the information set to which that node belongs. The branching factor of the tree is supposed to represent the number of choices available to each player at each node. In order to make things interesting, p is also assumed to be *non-injective*, thereby ensuring that at least one player gets to move more than once.

We shall set the following constraints on I , namely: (i) the information sets it assigns are exhaustive and pairwise disjoint; (ii) for each node $x \in G$, it holds that $x \in I(x)$; and (iii) nodes belonging to the same information set are assigned to the same player. Payoffs at the terminal nodes (leaves) of the tree are represented by *real-valued vectors*, whose i -th projections (for $0 < i \leq k$) represent the payoff for player i at that leaf.

However, there is nothing conceptual to gain in representing such generality, while there is much to lose in notational perspicuity. All the points that we want to make can be made equally well for a restricted class of games. Consequently, we make the following simplifying assumptions. We will assume only two players that move in a pre-determined order (with a player possibly moving more than once in a row). Accordingly, payoffs at the leaves are represented by *pairs* of real values. It will be convenient to introduce a function $q : G - \{a\} \rightarrow \{1, 2\}$ (where a is the root of G), that for each node x other than the root gives the player that moves at the *previous* node.

We will also restrict ourselves to games represented by *balanced binary trees*, i.e., games in which each player has precisely two choices at each node and all branches have the same length. Conventionally, the two choices are referred to as “moving left” and “moving right.” The trees are assumed to be “balanced,” i.e., such that all branches have the same length: Any unbalanced tree can be turned into a balanced one by adding nodes that are redundant from the point of view of the game (because they all lead to the same payoff vector). Similarly, we want information sets to contain nodes that have the same distance from any leaf: Information sets can thus be rearranged so that they contain only nodes of the same level in the tree. This can be accomplished as follows: If $x' \in I(x)$ is a node having the lowest level (i.e., the greatest distance from the root) among nodes in $I(x)$, when re-balancing the tree we put in the same information set all the nodes of the same level as x' that descend from nodes in $I(x)$ in the original tree. As an example consider the equivalent descriptions of Figures 5 and 6.

CONVENTION Assume two players, 1 and 2, of whom player 1 moves first, so that the root of

Figure 6: The “balanced” version of Figure 5.

the tree represents a choice for 1. In what follows, a always denotes the *root* of G . Call a node *final* if it is non-terminal, but both its children are terminal. We write $x \sim y$ if x and y are “siblings,” i.e., they are immediate successors of the same node; we also say that two paths are “siblings” if their initial nodes are. For any node x we denote by x_r and x_l its right- and left-hand successors, respectively. Moreover, by a *path* we mean a possibly empty sequence of nodes, each one of which is the successor of the previous one and the last of which is a leaf. A maximal path is called a *branch*. If x is a leaf, $\pi(x)$ represents the associated payoff vector, and if s is a path of length k , we write $\pi(s)$ instead of $\pi((s)_k)$. Also, we write $\pi(x), \pi(y) > \pi(z)$ to mean $\pi(x) > \pi(z)$ and $\pi(y) > \pi(z)$. We use x, y, z as variables for nodes and s, t, u as variables for paths. If $b = \langle b_0, \dots, b_k \rangle$ is any sequence and $i = 0, \dots, k$, we set $(b)_i = b_i$. If s and t are sequences of nodes, their concatenation is denoted by $s * t$. For notational convenience, let

$$\text{dir}(s, s') \iff [(s)_1 = ((s)_0)_r \leftrightarrow (s')_1 = ((s')_0)_r].$$

Intuitively, $\text{dir}(s, s')$ tells us that the two paths s and s' go in the same direction (i.e., right or left). Whenever two paths s and s' have different directions, we write $\neg \text{dir}(s, s')$. In practice, we will use $\text{dir}(s, s')$ only when s and s' are sibling paths.

We now want to define a function π^* that associates with each information set the set of all paths that are (weakly) undominated at that information set for the player whose turn it is to move at that set. The definition of $\pi^*(I(x))$ is by recursion on the level of x . This is sound, since all the nodes in $I(x)$ are assumed to be of the same level.

Suppose x is a final node (i.e., a lowest non-terminal node). Suppose x belongs to an information set $I(x)$ at which player i is called upon to move, and let s be a path starting from a node in $I(x)$. We say that a node x (weakly) dominates a sibling node y for player j (where $j \neq i$) relative to a set P of paths, and write $\text{Dom}(x, y, j, P)$, if and only if $x \sim y$, and

$$\begin{aligned} & (\forall s \in P)(\forall s' \in P)[(s)_0 = x \wedge (s')_0 = y \rightarrow (\pi(s))_j \geq (\pi(s'))_j] \wedge \\ & (\exists s \in P)[(s)_0 = x \wedge (\forall s' \in P)(s' \neq s \wedge (s')_0 = y \rightarrow (\pi(s))_j > (\pi(s'))_j)]. \end{aligned}$$

Note that $\text{Dom}(x, y, q(x), P)$ represents the fact that x dominates y for player $q(x)$, who moves at the *previous* information set. In the definition below the parameter P is replaced by the set $S_x = \{s : (s)_0 \in I(x)\}$, i.e. the set of paths whose initial nodes belong to $I(x)$.

Let us now define the set of paths starting at undominated nodes of information set $I(x)$ as:

$$\text{ND}(I(x)) = \{s : \neg \exists y \text{Dom}(y, (s)_0, q(x), S_x)\}$$

Thus whenever a player is choosing at an information set containing final nodes, that player will first consider whether any node that belongs to his information set is weakly dominated for the previous player (relative to siblings paths beginning at that information set). In which case he restricts his attention to undominated nodes. He will then eliminate any paths starting from undominated nodes that correspond to a weakly dominated move (from his own point of view). Note that this procedure corresponds to applying condition (R): Eliminating nodes that are weakly dominated for the previous player is consistent with common knowledge of rationality and it embodies the essence of Forward Induction.

To say that a path s starting from a final node corresponds to a (weakly) dominated move for player i , relative to a set P of paths, we write:

$$\begin{aligned} \text{dom}(s, i, P) &\iff \\ (\forall z \in I((s)_0))(\forall t, t' \in P)[(t)_0 = (t')_0 = z \wedge \text{dir}(s, t) \wedge \neg \text{dir}(s, t') \\ &\quad \rightarrow (\pi(t))_i \leq (\pi(t'))_i] \wedge \\ (\exists z \in I((s)_0))(\forall t, t' \in P)[(t)_0 = (t')_0 = z \wedge \text{dir}(s, t) \wedge \neg \text{dir}(s, t') \\ &\quad \rightarrow (\pi(t))_i < (\pi(t'))_i]. \end{aligned}$$

Recall that when x is a final node, s is a path comprising only two nodes.

We now construct $\pi^*(I(x))$, where $I(x)$ is a final information set at which player i is called upon to play, as follows:

$$\pi^*(I(x)) = \{s : (s)_0 \in I(x) \wedge \neg \text{dom}(s, i, \text{ND}(I((s)_0))) \wedge s \in \text{ND}(I((s)_0))\}.$$

Whenever a final information set is a singleton, our procedure applies without change. Consider for instance Figure 1: node y is strictly dominated by sibling node x_1 for player 1. Upon reaching her information set, player 2 will thus only consider those paths that start from node y' .

Having defined π^* for all final information sets, we proceed to recursively define π^* for information sets whose nodes are not final. Let $I(x)$ be such a non-final information set. We now take the set of all paths in $\pi^*(I(y))$, where y is a child of a node in $I(x)$: by inductive hypothesis, $\pi^*(I(y))$ is defined. For notational convenience, let

$$T_x = \{(x') * s : (x' \in I(x) \wedge y \text{ is a child of } x') \wedge s \in \pi^*(I(y))\}.$$

Since our goal is to construct bottom up the undominated paths for the entire game, any undominated path starting at a node in $I(y)$ must be extended upward to its predecessor in $I(x)$. These are precisely the paths contained in T_x . Let ND be defined as before, except that T_x is used in place of S_x :

$$\text{ND}(I(x)) = \{s : \neg \exists y \text{ Dom}(y, (s)_0, q(x), T_x)\}$$

Then we can define:

$$\begin{aligned} \pi^*(I(x)) &= \\ \{s \in T_x : (s)_0 \in I(x) \wedge \neg \text{dom}(s, i, \text{ND}(I((s)_0))) \wedge s \in \text{ND}(I((s)_0))\}. \end{aligned}$$

Intuitively, T_x is the set of all paths s starting from a node x' in $I(x)$ that are obtained by extending a path t in $\pi^*(I(y))$ (where y is a child of some $x' \in I(x)$) to the node immediately above $(t)_0$. Similarly to the case of x final, the player who moves at $I(x)$ selects those paths that begin with a node that is not dominated by a sibling from the point a view of the previous player, and then, *from among these paths*, the player selects those paths that correspond to weakly undominated moves.

Let us see how this definition works in the game of Figure 1. Observe first of all that node y' dominates node y from the point of view of player 1, who moves at the common parent of y and y' . Therefore player 2, when considering whether she has an undominated move, only considers paths beginning at y' . When thus restricting her attention, player 2 clearly has exactly one undominated move, namely R . At node x , player 1 will then choose between move a and move c : since the former dominates the latter, our procedure identifies (a, R) as the unique solution.

The game of Figure 3 is more complex. To apply our procedure we have to transform the subgame starting at node w into an equivalent subgame, in which node z has been eliminated. In the new subgame player 1 has three possible moves at w : either he chooses $D3$, or he chooses $A3T$ or $A3B$. That is, player 1 may choose to play $2?3$, thus ending the game, or engage in a simultaneous "battle of the sexes" with player 2. When considering the nodes in $7(5)$, player 2 will realize that node s' is dominated by node xv for player 1. Player 2 will then restrict her attention to the undominated paths starting from node s ; $x^*(I(s))$ then returns the path (s, s') that corresponds to move L . It is now possible to compute $TT^*(U)$, which gives the path (w, s, s') corresponding to the combination of moves $(A^T L)$. In turn, $7r^*(I(y))$ gives the path (y, y') corresponding to the move D_2 , and finally $ic^*(I(x))$ returns the path (x, y, y') corresponding to the combination of moves $(A_i, 2?2)$. According to our procedure, the unique solution for the game of Figure 3 is the path corresponding to (A_i, L) .

The procedure $TT^*(U)$ identifies a set of paths as "solutions" to the game. Such solutions might or might not be Nash equilibria, and some Nash equilibria might not be in $TT^*(U)$ (e.g., because they contain weakly dominated strategies). But if $TT^*(U)$ returns a unique path, such a path must correspond to a Nash equilibrium.

Let us now consider the concept of Nash equilibrium for extensive form games. A branch through the tree corresponds not to one but to *two* strategies, one for each player. We thus have to compare the payoff of a given branch (for a given player) with the payoffs of all other branches that embody alternative strategies, keeping fixed, however, the elements of the original branch that correspond to a strategy for the other player. We now proceed to capture this formally as follows. By a *move* we understand a length-two sequence of nodes such that the second is a child of the first. A *strategy*, in turn, is a set of moves, some of which might never be played if the strategy is chosen.

We start from a path s and partition it into two sets $M_1(s)$ and $M_2(s)$ corresponding to the moves of the two players: M_1 is the set of all length-two paths (x, y) such that (x, y) is a subpath of s and player 1 moves at x . Similarly, M_2 is the set of all length-two paths (x', y') such that (x', y') is a subpath of s , and player 2 moves at x' . We then expand $M_i(s)$ to a set S that is (1) "move-uniform" and, moreover, (2) it contains a response to each possible move of player $3 - i$. To say that S is move-uniform means that it satisfies the following condition: if $(x, y) \in S$, $x' \in \text{left}(x)$ and y' is the left- or right-hand child of x' according as y is the left- or right-hand child of x , then $(x', y') \in S$ too. To say that S contains a response to each possible move of $3 - i$ means that it satisfies the following condition: if $(x, y) \in S$ and $3 - i$ moves at y and z is a child of y , then there is exactly one move (z, u) that belongs to S . It is clear that $M_i(s)$ can always be extended (non-deterministically) to a set S satisfying the above two conditions. S corresponds to a strategy in strategic form games. Let us define a set of moves S to be a *strategy* for player i if it is a minimal set of moves containing $M_i(s)$ for some branch s and satisfying conditions (1) and (2).

Given a strategy S_1 for player 1 and a strategy S_2 for player 2, there is at most one branch s such that all and only its length-two segments are contained in $S_1 \cap S_2$, in which case (by slightly abusing the language) we will write $s = S_1 \cap S_2$. Define $u_1(S_1, S_2)$, the payoff for player 1 of playing strategy S_1 against strategy S_2 , as $u_1(s)$ if there is a unique branch s contained in $S_1 \cap S_2$, and set $u_1(S_1, S_2) = -\infty$ otherwise; u_2 is similarly defined for player 2).

A pair of strategies (S_1, S_2) (for players 1 and 2, respectively) is a Nash equilibrium if and only if:

- for any strategy S' for 1, $u_1(S', S_2) \leq u_1(S_1, S_2)$; and
- for any strategy S' for 2, $u_2(S_1, S') \leq u_2(S_1, S_2)$.

THEOREM 3.1 Let $ir^*(I(a)) = \{s\}$, where a is the root of the tree, and let S_1, S_2 be two strategies such that $s = S_1 \cap S_2$. Then (S_1, S_2) is a Nash equilibrium.

Proof First of all notice that $ir^*(I(a))$ is never empty (as can be easily shown by induction on the height of a), so our hypothesis comes down to the fact that $ir^*(I(a))$ contains at most one path. Now suppose for contradiction that (S_1, S_2) were not a Nash equilibrium. Then one of S_1, S_2 is dominated by some strategy S' (from the point of view of player i) in the sense that, for example, $u_i(S', S_2) > u_i(S_1, S_2)$. Suppose S_1 is dominated in this sense by S' for player 1. Of course, this can only happen if the two strategies S', S_2 intersect to give a branch, otherwise $u_i(S', S_2) = -\infty$. Let $s' = S' \cap S_2$; it follows, in particular, that s and s' must represent the same unique sequence of moves for player 2. Our hypothesis gives $u_i(s', s) > u_i(s, s)$. Furthermore, notice that this last fact implies that no move along s' is ever weakly dominated by a move along s for player 1.

Say that two paths t and t' are *undivided* at an information set $I(x)$ if and only if either they are both in $ir^*(I(x))$ or they are both outside of $ir^*(I(x))$. Observe that s and s' must intersect the same information set (if anything, they must meet at the root). Consider the lowest information set $I(x)$ intersected by both s and s' ; let k be the height of x . It is clear from our procedure (and can be shown by induction on the height of x) that if s and s' are undivided at such a lowest $I(x)$ then they are always undivided, in the sense that they are undivided at any node y not below x . In what follows, we consider two cases according as player 1 or player 2 moves at $I(x)$.

Suppose player 2 moves at $I(x)$. Then $(s) \wedge (s')$, otherwise (since $I(x)$ is the lowest information set at which s and s' intersect) s and s' would represent *different* moves for player 2. This is impossible, since we are keeping player 2's strategy S_2 fixed. By the same token, we must have $u_i(s, s')$. Then the only way for s and s' to be divided at $I(x)$ is if $(s)k$ dominates $(s')k$ from the point of view of the previous player: such a player cannot be player 2 (or else s and s' would again represent different moves for 2), and it cannot be player 1 (by hypothesis). Then, s and s' are undivided at $I(x)$, which implies that they are always undivided and in particular undivided at the root. That is, $s \in ir^*(I(a))$ as well, against our assumption.

Now suppose player 1 moves at $I(x)$. Observe that the hypothesis $u_i(s', s) > u_i(s, s)$ implies that $(s)k \wedge (s')k$. Otherwise, the only way in which s can survive the procedure ir^* while s' cannot is if moving in the direction of s dominates moving in the direction of s' for the player who moves at $(s)k$, which is precisely what the hypothesis rules out.

As before, it suffices to show that $s \in ir^*(I(x))$ (where, again, $I(x)$ is the lowest information set intersected by both s and s'). So suppose $s \notin ir^*(I(x))$. Now, the hypothesis $u_i(s', s) > u_i(s, s)$ implies that s' is ruled out for some reason other than the fact that moving in the direction of s dominates moving in the direction of s' for the player who moves at $I(x) = I((s)k)$ (i.e., player 1). It follows that s' can be ruled out only in one way. That is, s' can only be ruled out if $(s)k$ dominates $(s')k$ from the point of view of the previous player: but this cannot be player 1 (by hypothesis), and it cannot be player 2, for then s and s' would contain *different* moves for player 2, which, as we have seen, is impossible.

It is worth noting here that there is a similarity between our procedure and the iterated elimination of dominated strategies in strategic form games. In both cases, we have that if the outcome is a unique branch or, respectively, a pair of strategies, then it must be a Nash equilibrium. Indeed, the resemblance runs deeper. Our procedure can be thought of as a way of performing iterated elimination of dominated strategies in a given particular order. Such an order, however, is far from arbitrary, since it is dictated by the topology of the tree representing the game in extensive form. That is, the particular order is obtained from information which is lost in the strategic form. Moreover, it is precisely the use of such information that allows us to cast backwards and forward induction in a unified framework.

We now proceed to give a first-order theory T that, employing function π^* , allows us to predict the players' behavior. Since the definition of π^* is formalizable in Primitive Recursive Arithmetic, we shall assume that T contains enough arithmetic to carry out that definition. Moreover, T will have to contain axioms describing the tree representing the game and the structure of the payoffs at the leaves. We shall assume that all these "structural" axioms have been specified, and proceed to give the "behavioral" axioms.

First of all, we want to say that if a certain non-terminal node is reached, then the player whose turn it is to move will choose exactly one of the possible moves. We introduce predicates $L(x)$ and $R(x)$ with the intended meaning that the player (whose turn it is to move) moves left and, respectively, right at node x . If a is the root of the tree, we introduce the axiom

$$R(a) \leftrightarrow \neg L(a). \quad (1)$$

Moreover, for each non-terminal node x other than the root, we proceed as follows: let $x_1, \dots, x_{n+1} = x$ be the nodes on the path from the root to x , and let Q_i be R or L according as x_{i+1} is reached by moving left or right at x_i . Then we introduce an axiom saying that

$$Q_1(x_1) \wedge \dots \wedge Q_n(x_n) \rightarrow (R(x) \leftrightarrow \neg L(x)). \quad (2)$$

Next, we introduce an individual constant s representing (a suitable coding) of the set of undominated paths. We introduce an axiom to the effect that such a set is obtained by our procedure:

$$s = \pi^*(I(a)), \quad (3)$$

where, again, a is the root of the tree. Finally, we introduce axioms saying that players only move along the undominated paths: let y be any non-terminal node and suppose it has height (= number of predecessors) k . Then we have the axioms:

$$L(y) \rightarrow \exists s \in s ((s)_k = y_l); \quad (4)$$

$$R(y) \rightarrow \exists s \in s ((s)_k = y_r). \quad (5)$$

This completes our specification of T , which will then comprise (1)–(5) as behavioral axioms.

4 Towards a Complete Theory

We now indicate how the previous theory T can be modified to handle real deviations (i.e., those deviations that involve (weakly) dominated strategies). We are going to interpret T as a *default theory* $T' = (W, D)$, where W is a set of first-order axioms and D a set of normal defaults. A normal default is a weak inference rule of the form $\varphi \rightsquigarrow \psi$, interpreted as saying "if φ is known, and ψ is consistent, infer ψ " The sense in which ψ has to be consistent in order for it to be inferable is made precise in Default Logic (see Reiter [10] for details).

We have to specify W and D . In our theory, W comprises all the "structural axiom," i.e., whatever arithmetic is necessary to describe the game and compute π^* , along with a suitable coding of the game and associated payoffs. As before, we will leave this unspecified. Moreover, W will contain all formulas of the form (1) and (2).

On the other hand, D will specify the set of paths to be used in inferring the players' behavior. This set of paths will vary according as the node that has been reached in the game lies on or off the undominated paths.

Let a be the root, and T a propositional constant representing truth. Then we have a default to the effect that at the beginning of the game we use the undominated paths provided by $\neg r^*(J(a))$:

$$T \wedge s = \neg r^*(J(a)); \quad (6)$$

for any other node x , let $x_1, \dots, x_{n+1} = x$ be the nodes on the path from the root to x , and let Q_i be R or L according as x_{t+1} is reached by moving left or right at X_t . Then we introduce a default of the form:

$$Q_1(x_1) \wedge \dots \wedge Q_n(x_n) \wedge s = \neg r^*(J(x)). \quad (7)$$

This completes our specification of T' . This theory is *complete* in the sense that it can be augmented with information to the effect that a real deviation has taken place without becoming inconsistent. Moreover, it still allows us to say something about the game *after* a deviation, in the sense that it will still have an extension (in the sense of Reiter [10]) according to which all moves following a deviation still take place along one of the paths that are undominated in the subgame whose root is represented by the node at which the deviation has taken place. In this sense, T' embodies a principle of *local rationality* not dissimilar to the one presented in Bicchieri & Antonelli [5].

References

- [1] R. Aumann, *Backwards Induction and Common Knowledge of Rationality*, mimeo, University of Jerusalem, 1993.
- [2] C. Bicchieri, *Self Refuting Theories of Strategic Interaction: A Paradox of Common Knowledge*, **Erkenntnis** 30 (1989), pp. 69-85.
- [3] C. Bicchieri, *Knowledge-dependent Games: Backward Induction*, in Bicchieri & Dalla Chiara, **Knowledge, Belief, and Strategic Interaction**, Cambridge University Press, Cambridge 1992.
- [4] C. Bicchieri, **Rationality and Coordination**, Cambridge University Press, Cambridge, 1993.
- [5] C. Bicchieri & G.A. Antonelli, *Game-Theoretic Axioms for Bounded Rationality and Local Knowledge*, paper given at the Nobel Symposium on Game Theory, Björkborn, Sweden, June 1993.
- [6] R. Fagin (ed.), **Theoretical Aspects of Reasoning about Knowledge. Proceedings of the Fifth Conference (TARK 94)**, Morgan Kaufmann, San Francisco 1994.
- [7] E. Kohlberg & J.F. Mertens, *On the Strategic Stability of Equilibria*, *Econometrica* 54 (1986), pp. 1003-37.
- [8] D. Kreps and R. Wilson, *Sequential Equilibria*, **Econometrica** 50 (1982), pp. 863-94.
- [9] R.B. Myerson, *Credible Negotiation Statements and Coherent Plans*, **Journal of Economic Theory**, v. 48 (1989), pp. 264-303.
- [10] R. Reiter, *A Logic for Default Reasoning*, **Artificial Intelligence**, v. 13 (1980), pp. 81-132.
- [11] R. Selten, *Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit*, *Zeitschrift für die gesamte Staatswissenschaft* 121 (1965), 667-89.

- [12] R. Selten, *Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games*, *International Journal of Game Theory*, v. 4 (1975), pp. 25-55.
- [13] R. Stalnaker, *Knowledge, Belief, and Counterfactual Reasoning in Games*, forthcoming in the proceedings of the second Castiglioncello conference 1992, edited by C. Bicchieri and B. Skyrms.





Figure 1

	L	R
c	3,3	3,3
b	2,4	0,0
a	0,0	5,1

Figure 2



8
1980

D_2	A_1	A_2	A_3
rH rH	rH rH^{sk}	rH rH	rH rH
CM	ro	CO CO	CO CO
CM CM	rH $os.$	$0,0$	$0,0$
CM CM	$0,0$	rH	rH

es

$A_1 \approx$

$A_1 A_2$

$A_1 A_3 B$

Fig. 4




Figure 2




Figure 3