

Game Solution, Epistemic Dynamics and Fixed-Point Logics

Johan van Benthem^{*†}, Amélie Gheerbrant[†]

ILLC, University of Amsterdam, The Netherlands

johan@science.uva.nl, A.Gheerbrant@uva.nl

Abstract. Current methods for solving games embody a form of “procedural rationality” that invites logical analysis in its own right. This paper is a brief case study of Backward Induction for extensive games, replacing earlier static logical definitions by stepwise dynamic ones. We consider a number of analysis from recent years that look different conceptually, and find that they are all mathematically equivalent. This shows how an abstract logical perspective can bring out basic invariant structure in games. We then generalize this to an exploration of fixed-point logics on finite trees that best fit game-theoretic equilibria. We end with some open questions that suggest a broader program for merging current computational logics with notions and results from game theory. This paper is largely a program for opening up an area: an extended version of the technical results will be found in the forthcoming dissertation [26].

Keywords: Backward Induction, Extensive Games, Dynamic-Epistemic Logic, Fixed-Point Logic.

1. Game solution as rational procedure

Logic and games form a natural combination. On the one hand, there are “logic games” that analyze basic notions such as truth, proof, or model comparison, while on the other hand, standard logical systems have proved applicable to many basic issues in the foundations of game theory (cf. [5], [37]). This paper will concentrate on the second aspect.

*We thank Alexandru Baltag and Jonathan Zvesper for discussions at an earlier stage of this project. We also thank Tomohiro Hoshi, Jonathan Zvesper, and especially, Balder ten Cate and Cédric Dégremon for their comments on the present manuscript. We finally thank two anonymous reviewers of *Fundamenta Informaticae*.

[†]Address for correspondence: Institute for Logic, Language and Computation, University of Amsterdam, P.O. Box 94242, 1090 GE Amsterdam, The Netherlands

Logics that describe games In recent years, many logical analyses have been given of both strategic and extensive games, through introducing formal languages that describe game structure while raising logical questions of definability and axiomatization ([17], [4], [19], [29]). A benchmark for logics in this tradition has been the definition of *Backward Induction* (“BI” for short), the most common method for solving finite extensive games of perfect information ([39], [40]). In this same arena, basic foundational results have been obtained in epistemic game theory, endowing bare games with epistemic assumptions about players. A pilot result was the characterization of the BI outcome in terms of assuming common knowledge, or common true belief, in rationality, meaning that players choose those actions that they believe to be best for themselves ([1]).

Analyzing solution procedures Recently, [9] has suggested that the main focus here should be shifted: away from a static assumption of known or believed rationality to the underlying “procedural rationality” of *plausible procedures* that players engage in when analyzing and playing a game, and the way these result in stable limit models where rationality becomes common knowledge.¹ Thus, [4] shows how game-theoretic equilibrium fits with the computational perspective of fixed-point logics, and [11] gives several dynamic procedures that analyze BI. This paper will analyze these proposals further, and find their common mathematical background. This will then be our starting point for suggesting a more general line of investigation.

Basics of extensive games We assume some basic game theory, and we will work with *finite extensive games* of perfect information, i.e., finite trees with labelled nodes, where each node is either an end node, or an intermediate node that represents the turn of a unique player.² We will mostly think of 2-player games, though much of what we say generalizes to more players. While game trees with moves are simple computational structures, the essence of rational action arises with the way players evaluate outcomes. Thus, there is also a further *preference relation* for each player between end nodes (encoding complete histories) that we will take to be a total order in this paper, though this requirement could be generalized. Equivalently, such total evaluation orders may be represented in the form of numerical utility values for players at end nodes.

Backward induction We now define our basic procedure in a bit more detail:

Definition 1.1. (BI procedure for “generic” extensive games)

We call a game *generic* when, for each player, distinct end nodes have different utility values. On such games *Backward Induction* is this inductive algorithm:

“At end nodes, players already have their values marked. At further nodes, once all daughters are marked, the player to move gets her maximal value that occurs on a daughter, while the other, non-active player gets his value on that maximal node.”

A *strategy* for a player is a map that selects one move at each turn for that player. It is easy to see that BI generates a strategy for each player at her turns: go to the successor node that has your highest

¹Note that even the common word “solution” has an ambiguity between a procedure (“Solution is not easy”) and a static product of such a procedure (“Show me your solution”).

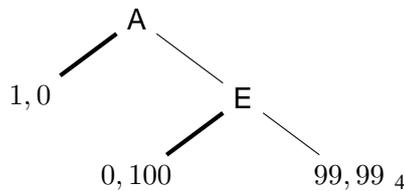
²Only towards the end, we will briefly consider games with imperfect information.

value. The resulting set of strategies is the “BI outcome”, that leads to a unique play of the game. We will call the set-theoretic union of all these strategies (still a function on nodes) *bi*. The BI procedure seems obvious, telling us players’ best course of action. And yet, it is packed with assumptions about how players behave that are worth high-lighting. For now, just note that the algorithm subtly changes its interpretation of values on the way. At leaves, these values encode plain utilities or preferences, but at nodes higher up in the game tree, the BI values clearly mix in additional considerations of plausibility, incorporating beliefs about what others will do.

Delicate cases BI can produce debatable outcomes, as in the next illustration:

Example 1.1. (A simple BI outcome)

In the following game, players’ preferences are encoded in utility values, as pairs (value for A, value for E). Backward Induction tells player E to turn left at her turn, which gives A a belief that this will happen, and so, based on this belief about his counter-player, A should turn left at the start, making both worse off than they might have been:³



The fact that the BI prediction or recommendation is not always intuitive has motivated much logical analysis of the procedure and the reasoning underpinning it. In this paper, we will not enter this debate. We neither endorse nor reject Backward Induction, but we merely take it as our point of entry into the logic of game solution procedures. Our starting point are three different proposals for explaining what makes BI tick, that we will explain in due course. But before getting there, let us first make a generalization of what we mean by Backward Induction.

2. From functional to relational strategies

Strategies as subrelations of the *move* relation A game-theoretic strategy is usually taken to be a function on nodes in a game tree, yielding a unique recommendation for play there. But in many settings, it makes more sense to think of strategies as nondeterministic *binary subrelations* of the total relation *move* (the union of all labelled actions in the game) that merely constrain further moves by selecting one or more as admissible. This is in line with the colloquial use of the term “strategy”, it also reflects a common view of plans for action, and technically, it facilitates logical definitions of strategies in propositional dynamic logic [10].

³People sometimes defend this outcome by saying that the given game is “competitive” – but that amounts to giving information about the players that is not explicit in the game tree. If such extra information is relevant to solution, we may need a richer notion of game from the start. We will return to this issue in Section 10 below.

⁴Frankly, we have dramatized things a bit here to catch the reader’s attention. Since the numbers just encode ordinal preferences, the same point might have been made with values 0, 1, 2 and 3. But the undesirable point remains that the computed outcome is not Pareto-optimal.

Relational BI, first version Indeed, one common numerical formulation of BI already has this relational flavor. We now drop the assumption that games are generic:

Definition 2.1. (Relational Backward Induction, first version)

Starting from the leaves, one now assigns values for players at nodes using the rule:

Suppose that E is to move at a node, and all values for daughters are known. The E-value is the maximum of all the E-values on the daughters, while the A-value is the minimum of the A-values at all E-best daughters.⁵

The relation *bi* arising from this algorithm connects nodes to all daughters with maximal values for the active player, of which there may be more than one. This method focuses on minimal values that can be guaranteed when doing the best within one's power.

Solution algorithms make assumptions about players But while this looks like an obvious numerical rule, it does embody special assumptions about players. In particular, taking the minimum value is a worst-case assumption that my counter-player does not care about my interests after her own are satisfied. But we might also assume that she does, choosing among her maximum nodes one that is best for me. In that case, the second numerical value in the algorithm would be a maximum rather than a minimum. And other options are possible.⁶ This variety of relational versions of game solution is not a problem. It rather highlights an important feature of game theory: mathematical "solution methods" are not neutral, they encode significant assumptions about players. But the variety does suggest that we first find a general base version of BI that is not too specific:

A minimal notion of rationality: avoid stupid moves Here is one logical analysis of the variety for relational versions of BI. Let us first view matters from a somewhat higher standpoint. Suppose that I need to compare different moves of mine, each of which, given the relational nature of the procedure, still allows for many leaves (end nodes) that can be reached via further *bi-play*.⁷ A minimal notion of *Rationality* would then say that

I do not play a move when I have another move whose outcomes I prefer.

A source of variety: different set preferences This seems plausible, but what notion of preference is involved here? It is easy to see that, in the above first version of the BI algorithm, the following choice is made. Player *i* preferred a set *Y* of leaves reachable by further *bi-play* to another set *X* if the *minimum* of its values for *i* is higher. That is, we have the following $\forall\exists$ pattern for set preference:⁸

$$\forall y \in Y \exists x \in X : x <_i y$$

⁵The dual calculation for values at A's turns is completely analogous.

⁶Of course, one might view such alternatives as calling for a change in players' utilities. We will not get into this perennial issue of game preference transformations here.

⁷In this perspective with total outcomes of the game, we make a shift from the original version of the BI algorithm, which looked at daughters of the current node only.

⁸Given that we have finite total orders, we could also replace this by

$$\exists x \in X \forall y \in Y : x <_i y$$

But clearly, staying with the same over-all notion of Rationality, there are several alternatives for comparisons between reachable sets of outcomes. One common notion of preference for Y over X in the logical literature ([43], [34]) is the $\forall\forall$ stipulation that

$$\forall y \in Y \forall x \in X : x <_i y$$

Relational backward induction, second version Clearly, avoiding moves that shouldn't be taken under this stronger preference is a weaker constraint on behavior of players. Still, it fits with a minimal game-theoretic solution procedure for strategic games called eliminating strictly dominated strategies ([35]). We will take this second relational version of Backward Induction as our running example:

Definition 2.2. (Relational Backward Induction, second version)

First, mark all moves as “active”. Call a move to a node x *dominated* if x has a sibling from which all reachable endpoints via active moves are preferred by the current player to all reachable endpoints via active moves from x itself. The second version of the BI algorithm works in stages:

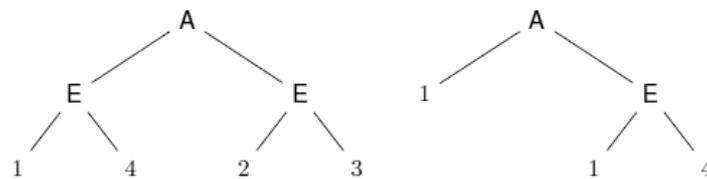
At each stage, it marks dominated moves in the $\forall\forall$ sense of set preference as “passive”, leaving all others active. In this preference comparison between sets of outcomes, the “reachable endpoints” by an active move are all those that can be reached via a sequence of moves that are still active at this stage.

In another well-known terminology, this says that players play a “best response”.

Henceforth, we will use BI to refer to this algorithm, and the subrelation of the total *move* relation produced by it at the end. It is a cautious notion of game solution making fewer assumptions about the behavior of other agents than the earlier version. Of course, the two versions agree on generic games, for which the subset of the move relation obtained as output is always a function.

Example 2.1. (Some comparisons)

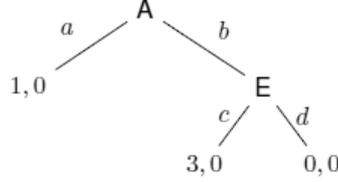
Consider the following two games, where the values indicated are utilities for player A. For simplicity, we assume that player E has no preference between her moves:



In the game to the left, our first version of Backward Induction makes A go right, since the minimum 2 is greater than the 1 on the left. But our cautious BI will accept both moves for A, as no move strictly dominates the other.

Moreover, both versions will accept all moves in the game to the right. This may seem strange, since most players would probably go right at the start: they have nothing to lose, and a lot to gain. But analyzing all variants for preference comparisons between sets of outcomes is not our focus here. We will return to the issue of further possible solution concepts in later sections.

Important remark Our style of analysis chooses one particular line toward generalizing Backward Induction to non-generic games. But others make sense, too, as pointed out by Cédric Dégrement. For instance, if one thinks of strategy profiles in Nash equilibrium, the following game would have two:



Both profiles (a, d) and (b, c) are in equilibrium. But our algorithm will leave both options for E, and tell A to go left. This paper will not address the alternative analysis of the BI-output in terms of sets of strategy profiles, leaving this as a challenge to fixed-point logics over richer models.

3. Defining BI as a unique static relation

Many definitions for the BI relation bi on generic games have been published by logicians and game-theorists (cf. the survey in [30]). Our point of departure here is a version involving a modal language of (a) labelled moves, i.e., binary transition relations a on nodes with matching modalities $\langle a \rangle$, plus (b) a modal preference operator interpreted as follows at nodes of a game tree:

$\langle pref_i \rangle \phi$: player i prefers some node where ϕ holds to the current one

The original result Here is a result from [13]:

Theorem 3.1. On generic games, the BI strategy is the unique function σ satisfying the following modal axiom for all propositions p - viewed as sets of nodes - for all players i :

$$(\text{turn}_i \wedge \langle \sigma^* \rangle (\text{end} \wedge p)) \rightarrow [\text{move}] \langle \sigma^* \rangle (\text{end} \wedge \langle pref_i \rangle p)$$

For a proof (a laborious but straightforward induction on finite tree depth), we refer to the cited paper. Here we just concentrate on the meaning of the crucial axiom, that may be brought out by a standard modal frame correspondence, where frame truth quantifies universally over all sets of objects for proposition letters ([16]). The frames here are games extended with one more binary relation σ . What we find is a notion of *Rationality* like before:

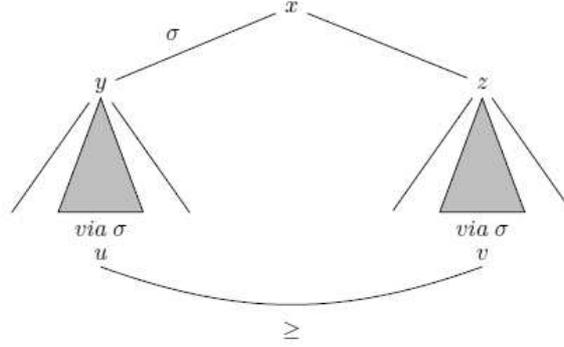
Fact 3.1. An extended game makes $(\text{turn}_i \wedge \langle \sigma^* \rangle (\text{end} \wedge p)) \rightarrow [\text{move}] \langle \sigma^* \rangle (\text{end} \wedge \langle pref_i \rangle p)$ true for all i at all nodes iff it has this property for all i :

RAT-1: No other available move for the current player i yields a set of outcomes by further play using σ that has a higher minimal value for i than the outcomes of playing σ all the way down the tree from the current node.

Proof:

This is a standard modal correspondence argument that we omit. The correspondence language uses the reflexive-transitive closure of the relation σ , but this is a simple extension of known techniques ([7]). \square

The typical picture to keep in mind here, and also later on, is this:



RAT-1 is equivalent to this *confluence property* for action and preference:

$$\text{CF1} : \bigwedge_i \forall x (\text{turn}_i(x) \rightarrow \forall y (\sigma(x, y) \rightarrow (\text{move}(x, y) \wedge \forall u ((\text{end}(u) \wedge \sigma^*(y, u)) \rightarrow \forall z (\text{move}(x, z) \rightarrow \exists v (\text{end}(v) \wedge \sigma^*(z, v) \wedge v \leq_i u))))))9$$

This $\forall\forall\forall\exists$ form is a comparison between sets of outcomes that negates an earlier notion of preference: the *minimum value* on the reachable endpoints after z is not larger than that after y . It is easy to show that any relation σ with this property matches the **BI** solution level by level on generic games.

Capturing BI in logical terms But now let us look at our favored relational generalization of **BI**. First, we reformulate the stated non-dominance property:

RAT-2. No alternative move for the current player i guarantees outcomes via further play using σ that are all strictly better for i than all outcomes resulting from starting at the current move and then playing σ all the way down the tree.

A logical formula defining this has the following $\forall\forall\exists\exists$ form:

$$\text{CF2} : \bigwedge_i \forall x \forall y ((\text{turn}_i(x) \wedge \sigma(x, y)) \rightarrow (\text{move}(x, y) \wedge \forall z (\text{move}(x, z) \rightarrow \exists u \exists v (\text{end}(u) \wedge \text{end}(v) \wedge \sigma^*(y, v) \wedge \sigma^*(z, u) \wedge u \leq_i v))))10$$

Theorem 3.2. **BI** is the largest subrelation of the *move* relation in a finite game tree satisfying the two properties that (a) the relation has a successor at each intermediate node, and (b) **CF2** holds.¹⁰

Proof:

First, the given algorithm clearly leaves at least one active move at each node, by the definition of preference. Moreover, at the final state, when no more deactivations occur, **CF2** must hold: there are no more dominated moves, and that is what it says.

⁹One could change the formal language for **CF1** here to a more technical first-order one avoiding the closure operator – but for our main points, such variations are not important.

¹⁰We say “largest” in this formulation because in the presence of more than one best successor, different subrelations of the *move* relation might satisfy **CF2**. Note that there need not be a largest relation satisfying a given structural property, but in this particular case, it does.

That the relation defined in this way is maximal may be seen as follows. If we reactivate anywhere a move that is inactive, that move had disappeared at some stage because it was dominated there by another move. But then it would still be dominated in the whole tree by the same move. For, all that can have happened in the further stages of the algorithm is that fewer endpoints have become reachable through active paths from the two moves, and their $\forall\forall$ -dominance relationship then persists.

Conversely, if we have any subrelation of the *move* relation with the given two properties, it is easy to see by induction on the depth of subtrees that all its moves survive each stage of the above main BI procedure, by the definition of the elimination step. \square

We now make these same points about the procedure more syntactically, by inspecting the syntax of CF2. We can restate this in terms of a well-known formalism: the first-order fixed-point logic LFP(FO) (cf. [24]):¹¹

Theorem 3.3. The BI relation is definable in LFP(FO).

Proof:

Indeed, the definition involves just one *greatest fixed-point* in addition to the transitive closure operations. This fixed-point is in the language of LFP(FO): all occurrences of the predicate symbol S in the relevant formula are positive:

$$\begin{aligned} \text{BI}(x, y) = \nu S, xy. (& \text{move}(x, y) \wedge \bigwedge_i (\text{turn}_i(x) \rightarrow \forall z (\text{move}(x, z) \\ & \rightarrow \exists u \exists v (\text{end}(u) \wedge \text{end}(v) \wedge S^*(y, v) \wedge S^*(z, u) \wedge u \leq_i v))) \end{aligned})^{12} \quad \square$$

This definition will be our point of reference in what follows. Interestingly, it is both a static description of the BI relation and also a definition of a procedure computing it. For, we can now use the standard defining sequence for a greatest fixed-point, starting from the total *move* relation, and see that its successive decreasing approximation stages S^k are exactly the “active move stages” of the above algorithm. We will refer to these stages S^k at several places in what follows. In our view, fixed-point logics are attractive since they analyze both the statics and dynamics of game solution.

In the following sections, we extend this theme by looking at two further logical ways of construing the Backward Induction procedure that have been proposed in recent years.

4. A dynamic-epistemic scenario: iterated announcement of rationality

Here is another procedural line on Backward Induction as a rational process. [9] proposed an analysis in the spirit of current *dynamic-epistemic logics* that describe acts of information flow, such as public announcements or observations ([23], [12]). The following analysis of BI takes it to be a process of prior off-line deliberation about a game by players whose minds proceed in harmony - though they need not communicate in reality.¹³

¹¹In terms of [6], the syntax of CF2 has dual “PIA form”, guaranteeing that the union of all relations satisfying CF2 exists, while a small extra argument gives the existence.

¹²We can also replace the reflexive transitive closures S^* by definitions in LFP(FO).

¹³Compare also the dynamic agreement procedures studied in [25].

Solving games by announcements of rationality The following analysis uses the dynamic epistemic logic of public announcements $!\phi$ saying that some proposition ϕ is true. These transform a current epistemic model \mathfrak{M} into its submodel $\mathfrak{M}|\phi$ whose domain consists of just those worlds in \mathfrak{M} that satisfy ϕ . [9] makes the solution process of extensive games itself the focus of a PAL style analysis:

Definition 4.1. (Node rationality)

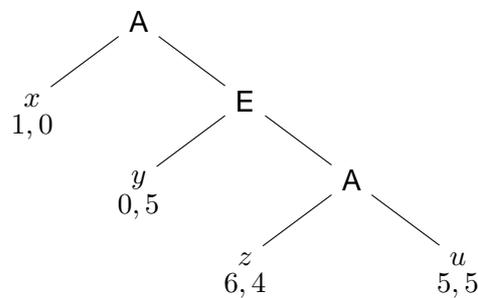
As before, at a turn for player i , a move to a node x is *dominated* by a move to a sibling y of x if every history through x ends worse, in terms of i 's preference, than every history through y . Now **rat** says that “at the current node, no player has chosen a strictly dominated move in the past coming here”.

This makes an assertion about nodes in a game tree, viz. that they did not arise through playing a dominated move. Some nodes will satisfy this, others may not. Note that we do not say that every node in the game satisfies **rat**: we merely say that it is an informative property of nodes. Thus, announcing this formula as a fact about the players of a game is informative, and it will in general make the current game tree smaller.

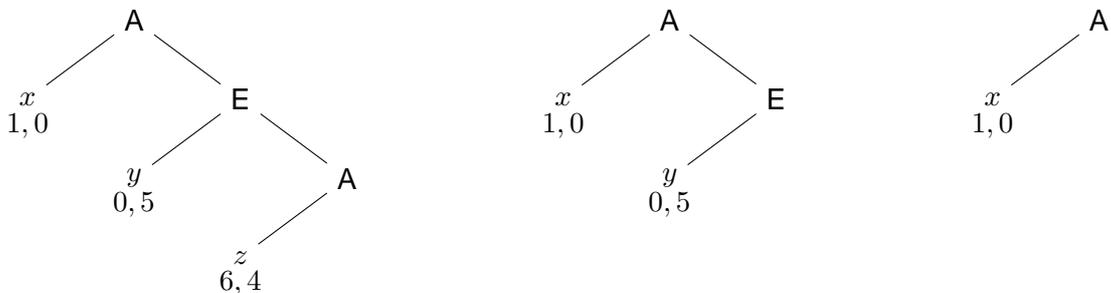
But then we get a dynamics as in famous puzzles like the Muddy Children, where *repeated assertions* of ignorance eventually produce enough information to solve the whole puzzle. In our case, in the new smaller game tree, new nodes may become dominated, and hence announcing **rat** *again* (saying that it *still* holds after this round of deliberation) makes sense, and so on. This process of iterated announcement must always reach a limit, that is, a smallest subgame where no node is dominated any more:

Example 4.1. (Solving games through iterated assertions of Rationality)

Consider a game with three turns, four branches, and pay-offs for A, E in that order:



Stage 0 of the procedure rules out point u (the only point where Rationality fails), Stage 1 rules out z and the node above it (the new points where Rationality fails), and Stage 2 rules out y and the node above it. In the remaining game, Rationality holds throughout:



In such generic games, the BI solution emerges step by step. [9] shows that the actual Backward Induction path for extensive games is obtained by repeated announcement of the assertion rat to its limit. We repeat some relevant notions from dynamic-epistemic logic:

Definition 4.2. (Announcement limit)

For each epistemic model \mathfrak{M} and each proposition ϕ that is true or false at points in the model, the announcement limit $(\phi, \mathfrak{M})^\#$ is the first model reached by successive announcements $\phi!$ that no longer changes after the last announcement is made.

That such a limit exists is clear for finite models, since the sequence of submodels is weakly decreasing.¹⁴ There are two possibilities for the limit model. Either it is non-empty, in which case ϕ holds in all nodes, meaning that it has become common knowledge (the *self-fulfilling* case), or it is empty, meaning that the negation $\neg\phi$ has become common knowledge (the *self-refuting* case). Both occur in concrete puzzles, though generally speaking, rationality assertions like rat tend to be self-fulfilling, while the ignorance statement that drives the Muddy Children is self-refuting: at the end, it holds nowhere.

Capturing BI by iterated announcements With general relational strategies, the iterated announcement scenario produces the earlier $\forall\forall\exists\exists$ version of Backward Induction:

Theorem 4.1. In any game tree \mathfrak{M} , $(!\text{rat}, \mathfrak{M})^\#$ is the actual subtree computed by BI.

This can be proved directly, but it also follows from our next observations. For a start, it turns out easier to change the definition of the driving assertion rat a bit. We now only demand that the current node was not arrived at directly via a dominated move for one of the players. This does not eliminate nodes further down, and indeed, announcing this repeatedly will make the game tree fall apart into a forest of disjoint subtrees – as is easily seen in the above examples. These record more information.

Sets of nodes as relations Here is an obvious fact about game trees. Each subrelation R of the total *move* relation has an obvious unique corresponding set of nodes $\text{reach}(R)$ consisting of the set-theoretic range of R plus the root of the tree (we add the latter for convenience). And vice versa, each set X of nodes induces a unique corresponding subrelation of the *move* relation $\text{rel}(X)$ consisting of all moves in the tree that end in X .

With this simple connection, we can link the earlier approximation stages BI^k for Backward Induction (i.e., the successive relations computed by our earlier procedure) and the stages of our public announcement procedure. They are in harmony all the way:

Fact 4.1. For each k , in each game model \mathfrak{M} , $\text{BI}^k = \text{rel}((!\text{rat})^k, \mathfrak{M})$.

Proof:

By induction on k . The base case is obvious: \mathfrak{M} is still the whole tree, and the relation BI^0 equals *move*. Next, consider the inductive step. If we announce rat again, we remove all points reached by a move that is dominated for at least one player. These are precisely the moves cancelled by the corresponding step of the BI algorithm. \square

¹⁴Announcement limits also exist in infinite models, taking intersections at limit ordinals.

It follows also that, for each stage k ,

$$reach(\mathbf{BI}^k) = ((!rat)^k, \mathfrak{M}).$$

Either way, we conclude that the earlier algorithmic fixed-point definition of the BI procedure and van Benthem’s iterated announcement procedure amount to the same thing.¹⁵

Thus, one might say that the deliberation scenario is just a way of “conversationalizing” the underlying mathematical fixed-point computation. Still, it is of interest in the following sense. Viewing a game tree as an epistemic model with nodes as worlds, we see how repeated announcement of Rationality eventually makes this property true throughout the remaining limit model: in this way, it has made itself into *common knowledge*.

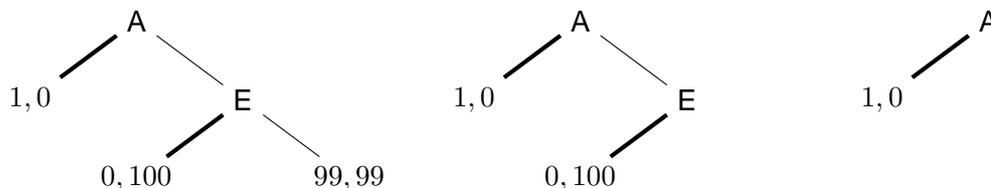
5. Another dynamic scenario: beliefs and iterated plausibility upgrade

Next, in addition to knowledge, consider the equally fundamental notion of *belief*. Many foundational studies in game theory (cf. the extensive discussion and references for belief-based game theory in [45]) view Rationality as choosing a best action *given what one believes* about the current and future behavior of the players. Indeed, this may be the most widely adopted view of game solution in the epistemic foundations of game theory today. We will first state a logical analysis of game solution in these terms, and then relate it to our earlier account of Backward Induction.

Backward Induction in a soft light An appealing take on the BI strategy in terms of beliefs uses “soft update” that does not eliminate worlds as above for announcements $!\phi$, but rearranges the plausibility order between worlds ([8]). A typical example is the *radical upgrade* $\uparrow \phi$ that makes all current ϕ -worlds best, and then puts all $\neg\phi$ -worlds underneath, while keeping the old ordering inside these two zones. Now recall our earlier observation that Backward Induction really creates *expectations* for players. All the essential information produced by the algorithm is then in the binary *plausibility relations* that it creates inductively for players among end nodes in the game, standing for complete histories. To see this, consider our running example once more:

Example 5.1. (The debatable BI outcome, hard and soft)

The hard scenario in terms of events $!rat$ removes nodes x from the tree that are reached via moves which are strictly dominated by moves to siblings of x as long as this can be done, resulting in the following sequence of stages:



¹⁵We leave the technical question open to which extent this is a more general technical method for switching between different types of predicate arities with fixed-points.

By contrast, a soft scenario does not remove nodes but modifies the plausibility relation. We start with all endpoints of the game tree incomparable (other versions would have them equiplausible). Next, at each stage, we compare sibling nodes, using this notion:

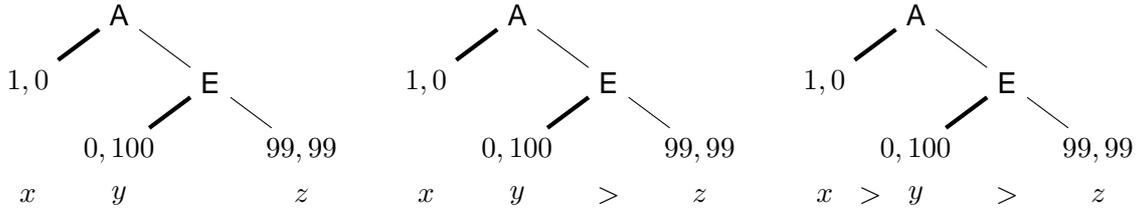
Definition 5.1. (Rationality in beliefs)

A move to a node x for player i *dominates* a move to a sibling y of x in beliefs if the most plausible end nodes reachable after x along any path in the whole game tree are all better for the active player than all the most plausible end nodes reachable in the game after y . Rationality* (rat^*) is the assertion that no player plays a move that is dominated in beliefs.

Now we perform a relation change that is like a radical upgrade $\uparrow \text{rat}^*$ ¹⁶:

If a move to a node x dominates a move to a sibling y of x in beliefs, we make all end nodes reachable from x more plausible than those reachable from y , keeping the old order inside these zones.

This changes the plausibility order, and hence the dominance pattern, so that belief statements can change their truth values – and a genuine iteration can start. Here are the stages for this procedure in the above example, where we use the letters x, y, z to stand for the end nodes or histories of the game:



In the first game tree, going right is not yet dominated in beliefs for A by going left. rat^* only has bite at E's turn, and an upgrade takes place that makes (0, 100) more plausible than (99, 99). After this upgrade, however, going right has now become dominated in beliefs, and a new upgrade takes place, making A's going left most plausible.

Here is a result stated without proof in [11]:

Theorem 5.1. On finite trees, the Backward Induction strategy is encoded in the final plausibility order for end nodes created by iterated radical upgrade with rationality in belief.

At the end of this procedure, players have acquired *common belief in rationality*. Let us now prove the result, using an idea from [3].

Strategies as special plausibility relations We first observe that each subrelation R of the total *move* relation induces a total plausibility order $\text{ord}(R)$ on leaves x, y of the tree.

Definition 5.2. (Leaf order from a sub-move relation)

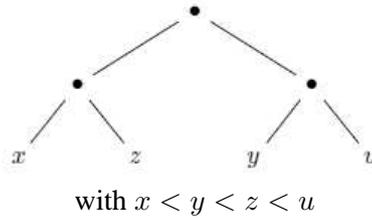
We put $x\text{ord}(R)y$ iff, looking upward at the first node z where the histories of x, y diverged, if x was reached via an R move from z , then so is y .

¹⁶We omit technical details here: plausibility upgrades may take place in subtrees, and hence one needs to work with submodels of the whole set of histories. Cf. [26] for details.

It is easy to see by inspection of trees that:

Fact 5.1. The relation $\text{ord}(R)$ is a total order on leaves.

Moreover, this total order on leaves is *tree-compatible*, meaning that, for any two leaves x, y , if z is the first splitting node above x, y as before, all leaves x' reached by taking the move toward x at z stand in the relation $\text{ord}(R)$ to all leaves y' reached by taking the move toward y . As an example, there can be no criss-crossing as in the following tree:



Definition 5.3. (Relational strategies from leaf order)

Conversely, any tree-compatible total order \leq on leaves induces a subrelation $\text{rel}(\leq)$ of the *move* relation, defined by selecting just those available moves at a node z that have the following property: their further available histories lead only to \leq -maximal leaves in the total set of leaves that are reachable from z .

Together, the two maps rel and ord give a precise meaning to the sense in which [3] can say that “strategies are the same as plausibility relations”.¹⁷

Now we can relate the computation in our upgrade scenario for belief and plausibility to the earlier relational algorithm for BI. Things turn out to be in harmony stage by stage:

Fact 5.2. For any game tree \mathfrak{M} and any k , $\text{rel}((\uparrow \text{rat}^*)^k, \mathfrak{M}) = \text{BI}^k$.

Proof:

The key point is as demonstrated in the earlier example of a stepwise BI solution procedure. When computing a next approximation for the BI-relation according to CF2, we drop those moves that are dominated by another available one. But this has the same effect as making the leaves reachable from dominated moves less plausible than those reachable from surviving moves. And that was precisely the earlier upgrade step. □

We conclude that the algorithmic analysis of Backward Induction and its procedural doxastic analysis in terms of forming beliefs amount to the same thing. Still, as with the iterated announcement scenario, the iterated upgrade scenario also has some interesting features of its own. One is that, for logicians, it yields fine-structure to the plausibility relations that are usually treated as primitives in models for doxastic logic. Thus games provide an underpinning for possible worlds semantics of belief that seems of interest per se.

¹⁷[44] relates our dynamic analysis to achieving the sufficient condition for the Backward Induction outcome given in [3].

Remark on syntax While we have stated the operations ord and rel purely semantically, they can obviously also be formulated as syntactic translations, and then the various logical definitions for Backward Induction given in this and preceding sections can also be directly transformed into each other. As before, we refer to [26] for details.

6. Midway conclusion: the stability of Backward Induction

Extensional equivalence, intensional difference We have now seen how three different approaches to analyzing Backward Induction turn out to amount to the same thing. To us, this means that the notion is stable, and that, in particular, its fixed-point definition can serve as a normal form. This motivates taking a closer look at fixed-point logics for game solution. Of course, as we have observed, extensionally equivalent definitions can still have interesting intensional differences in terms of what they suggest. For instance, we see the above analysis of strategy creation and plausibility change as one more concrete case study for a general conceptual issue: the fact that agents' beliefs and rational action are deeply entangled in the conceptual foundations of decision and game theory.

Dynamic instead of static foundations for game theory As we also said already, one key feature of our dynamic announcement and upgrade scenarios is this. In the terms of [9], they are *self-fulfilling*: ending in non-empty largest submodels where players have *common knowledge or common belief of rationality*.¹⁸ ¹⁹ Thus, this dynamic style of game analysis is a big change from the usual static characterizations of Backward Induction in the epistemic foundations of game theory. Common knowledge or belief of rationality is not assumed, but *produced* by the logic.

6.1. Test case: variants of Backward Induction

Are the preceding results just special effects for the notion of BI chosen here? As a test case, one can also look at the variant BI' considered in Section 2, where preference between sets of outcomes referred to ensuring a greater minimal value. This variant is studied in detail in the extended version of this paper (cf. [26]). We merely list some salient facts.

First, the $\forall\forall\forall\exists$ -type syntactic definition CF1 can no longer be used for an immediate fixed-point definition in LFP(FO). We would get

$$\begin{aligned} & \text{move}(x, y) \wedge \bigwedge_i (\text{turn}_i(x) \rightarrow \forall u ((\text{end}(u) \wedge S^*(y, u)) \\ & \rightarrow \forall z (\text{move}(x, z) \rightarrow \exists v (\text{end}(v) \wedge S^*(z, v) \wedge v \leq_i u))) \end{aligned}$$

where not all occurrences of the relation symbol S are positive. Still, we get

Theorem 6.1. The relational BI'-strategy is definable in the first-order inflationary fixed-point logic IFP(FO) using simultaneous fixed-points.

¹⁸We forego the issue of logical languages for explicitly *defining* the limit submodel.

¹⁹We also forego the further analysis of the limit behavior of upgrade actions on game models. For general models, [2] finds some curious phenomena, such as plausibility cycles, and they prove a general result stating when at least absolute beliefs stabilize in the limit.

For details on first-order and modal inflationary fixed-point logics, we refer to [24], [20].²⁰ ²¹

[26] also shows how the earlier iterated public announcement scenario for game solution and the iterated plausibility upgrade scenario both apply to Bl' , this time, with driving assertions derived from the condition CF1 . Thus we find the same stability in approaches that we saw for the earlier case of Bl .

7. Fixed-point logic of game solution: general issues

Which fixed-point logic? The preceding analysis raises the issue which of the many fixed-point logics known for computation best fit game solution methods. Candidates are the earlier-mentioned first-order logic with fixed-points $\text{LFP}(\text{FO})$ or the inflationary variant $\text{IFP}(\text{FO})$. These are general languages for recursion that work on any model. This generality is attractive, since we may want to investigate game solution procedures that are quite different from Backward Induction.

But an alternative take is also possible. One major feature of game solution procedures like Backward Induction is their exploiting the inductive structure of extensive games, via the *well-founded tree dominance order* toward the leaves.²² Such orderings allow for recursive definitions that yield uniqueness even without positive occurrence:

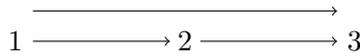
Example 7.1. (Fixed-points in modal provability logic ([18]))

On finite trees, any equivalence of the form $p \leftrightarrow \phi(p)$ where p occurs only “guarded” (that is, in the scope of at least one modality) in the formula ϕ , defines a unique proposition p . One proves this by induction on the well-founded tree order.

This includes examples that are not obviously in ordinary fixed-point logics:

Example 7.2. (Broader well-founded recursion)

Consider the definition $p \leftrightarrow \neg \Box p$. There is no obvious definition in the modal μ -calculus, and not even in its inflationary variant. Indeed, on a 3-nodes linear order



starting from any set as a value for p , this will stabilize with $p = \{2\}$. But it is easy to see that an inflationary bottom-up procedure for this formula stops in the pre-fixed point $\{1, 2\}$, and the deflationary top-down procedure stops in the post-fixed point \emptyset . Neither of these is even a fixed-point. What one can see more precisely in the straightforward approximation procedure, without forcing increasing or decreasing sets, is this: starting the iteration from any initial set will gradually get the predicate right, successively, at all nodes lying at increasing height from the leaves.²³

²⁰This is interesting, since [9] already observed how the limits of iterated public announcement procedures on modal models are definable in $\text{IFP}(\text{FO})$, and in fact, usually in the modal inflationary calculus, the extension of the modal μ -calculus by means of inflationary fixed-points ([20]).

²¹From the preceding fact, we can conclude (using [27], [33]) that there is an equivalent definition for Bl' in $\text{LFP}(\text{FO})$ after all, though the latter may involve extra predicates, with a computation no longer matching the natural stages of our algorithm.

²²But other recursions are possible, too. Both finite and infinite trees allow for recursive definitions over the well-founded tree order in the opposite *past direction* toward the root.

²³Balder ten Cate observes that the unique p with $p \leftrightarrow \neg \Box p$ is definable in the modal inflationary calculus ([20]) extended with simultaneous fixed-points, one for the ambient p , and one for the currently already stable subpredicate of p . We actually construct the formula defining the Bl' strategy (see Theorem 6.1) using the same idea.

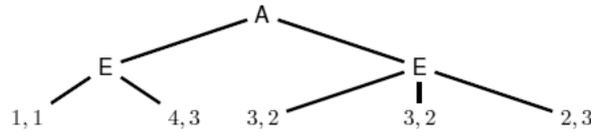
Of course, our analysis for Backward Induction did not use these simple modal languages, nor did it just use the simple tree dominance order. Still, it is easy to see by inspection of our earlier formulas and arguments that the following result holds

Fact 7.1. Stated as an equivalence, the Rationality principle CF2 defines a unique subrelation of the *move* relation by recursion on a well-founded order on finite trees: viz. the composition of the relations *sibling* and *dominance*.

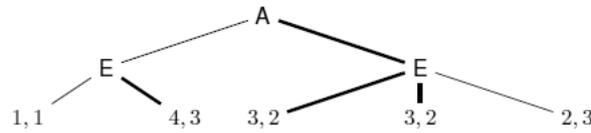
And this unique relation may also be computed for CF1 and other versions that lack positive syntax but do descend along the well-founded order. One can start with any subrelation of the *move* relation, and then compute according to the given instruction. At any stage k , the fixed-point relation stabilizes at points of distance $\leq k$ to the leaves.

Example 7.3. (Computing a fixed-point for CF1)

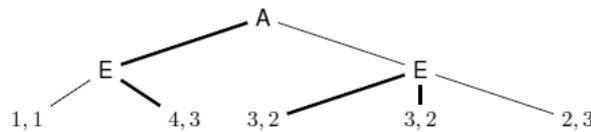
Consider this game, with values on leaves written as (value for E, value for A):



Let R^0 be the whole *move* relation. Then R^1 is marked in black below:



This still gets the fixed-point relation wrong at the root, but in the next stage we get the stable solution:



This suggests the introduction of a logical formalism for games that can access this order directly, generalizing modal syntax.²⁴ We will not do so here, but refer to [26] for a more abstract analysis of conditions under which such recursions are successful.

Fragments of fixed-point logics Of course, game solution procedures need not use the full power of logical languages that can define recursive procedures. Thus, there is a question *which fragments* are needed in our analysis. It might make sense to look at decidable fragments such as the *modal μ -calculus* - and indeed, [9] points out how the latter suffices, e.g., for defining the game solutions needed for Zermelo's Theorem. This may look too poor, since we often want to define relations on trees, and not

²⁴Relevant proposals in the literature include the "non monotone inductive definitions" of [22].

just unary predicates. But we have already seen how subsets of the *move* relation are encoded by unary predicates, so a lot can be done in this way. Still, the intriguing issue is this. Crucially, game solution intertwines two different relations on trees: the *move* relation and the *preference* relations for players on endpoints. And the question is what happens to the known properties of computational logics when we add such preference relations. In particular, the following intriguing issue then arises.

Potential problem: the complexity of rationality In logics of *action* and *knowledge*, it is well-known that apparently harmless assumptions such as Perfect Recall for agents make the bimodal logic undecidable, and sometimes even Π_1^1 -complete ([28]). The reason is that these assumptions generate commuting diagrams for actions *move* and epistemic uncertainty \sim satisfying a “confluence” property

$$\forall x \forall y ((\text{move}(x, y) \wedge y \sim z) \rightarrow \exists u (x \sim u \wedge \text{move}(u, z)))$$

that can serve as the basic grid cells in encodings of Tiling Problems in the logic. Thus, the logical theory of games with players that have perfect memory is more complex than that of forgetful agents ([14]).

But now consider the non-epistemic property of *Rationality* that mixes action and preference. The earlier properties CF1, CF2 have a similar flavor: they express the existence of a confluence diagram involving action and preference links. For instance, CF1 said this:

$$\begin{aligned} & \forall x \forall y ((\text{turn}_i(x) \wedge \sigma(x, y)) \rightarrow \forall z (\text{move}(x, z) \rightarrow \\ & \rightarrow \forall u ((\text{end}(u) \wedge \sigma^*(y, u)) \rightarrow \exists v (\text{end}(v) \wedge \sigma^*(z, v) \wedge v \leq_i u)))) \end{aligned}$$

So, what is the complexity of fixed-point logics for players with this kind of behavior? Can it be that Rationality, a widely used property meant to make behavior simple and predictable, actually makes its logical theory complex? Concrete instances of this open problem arise once we fix a sufficiently expressive logical language over trees: see below.²⁵

Language design and game equivalence As a final perspective, we mention that the choice of a best language for games is also correlated with the choice of an optimal notion of *structural equivalence* between games ([4]). The richer the equivalence, the stronger the language needed to capture its invariant properties. The options for languages that we have discussed here may also reflect the fact that there is no consensus yet on what such a structural notion of game equivalence should be.

8. The gist of it all: modal logics of best action

We have made a plea for analyzing game solution procedures explicitly in rich logics. This follows the program of *making strategies explicit* advocated in [10]. But while this is useful in some cases, there is also the opposite direction of judiciously hiding information about the machinery of strategies when it is not needed. In practical reasoning, we are often only interested in our *best actions* without all details of their justification. Game solution procedures take a model with actions and preferences, and

²⁵*Model-checking complexity and definability on finite trees.* Balder ten Cate has reminded us of the potential use of *descriptive complexity theory* ([31]) for studying finite games. First, checking for game solutions is related to model checking logical formulas, say, stating the intended effects of players’ strategies. As an example, since both LFP(FO) and IFP(FO) capture PTIME on finite models (given an enumeration order on the tree), it should be close to defining all “testable” properties of games. And other results in descriptive complexity theory may be game-theoretically relevant as well.

then compute a new relation of best action. As a mathematical abstraction, it would be good to extract a simple surface logic (a small modal fragment of complex fixed-point logics) for reasoning with best actions, while hiding most of the machinery:

Open problem Can we axiomatize the modal logic of finite game trees with a *move* relation and its transitive closure, turns and *preference* relations for players, and a new relation *best* as computed by Backward Induction?

We conjecture that we get a simple modal logic for the moves (these exist) plus a basic preference logic, while the modality $\langle best \rangle$ satisfies some obvious base laws plus one major bridge axiom that we already encountered earlier.²⁶

Fact 8.1. The following modal axiom corresponds to CF2 by standard techniques:

$$(\text{turn}_i \wedge \langle best \rangle [best^*](\text{end} \rightarrow p)) \rightarrow [\text{move}] \langle best^* \rangle (\text{end} \wedge \langle pref_i \rangle p)$$

In this concrete setting, the earlier problem returns that the Rationality assumption built into this logic may be a grid property leading to undecidability. Is the modal logic of best action decidable?

9. Further issues and extended game logics

There are many further lines for investigation following up on our stray observations. For instance, we want a more general view of possible representation languages, and on the notions of set preference that determine the dominance relation defining rationality. More generally, it would be of interest to connect our style of analysis for game solution more systematically with that found in epistemic game theory, where epistemic models are added describing what players know or believe about the course of the game.

In addition, some extensions to the games themselves seem natural:

Infinite games Can we extend our analysis to deal with *infinite games*? A transition to infinite ordinal sequences is easy to add to our iterated announcement or upgrade scenarios. Also, our general fixed-point definitions still make sense in this setting, though the special recursion over a well-founded tree dominance relation is no longer available. But there may be more to this generalization. Typically, in infinite trees, the reasoning changes direction, from “backward” to “forward”. Here is an illustration:

Example 9.1. (Weak Determinacy)

The following principle holds in all infinite game trees, for any condition ϕ on histories:

If player E has no strategy forcing ϕ at some stage s of the game, then A has a strategy for achieving a set of runs from s during all of which E never has a strategy forcing ϕ for the remaining game from then on.

In the notation of temporal game logics with *forcing modalities* $\{\}$, this says

$$\{E\}\phi \vee \{A\}G\neg\{E\}\phi$$

²⁶[36] has some related thoughts on “logics of solved games”.

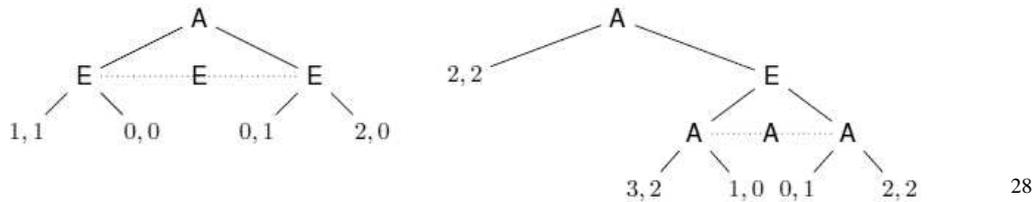
Here the reasoning is a typical inverse of Backward Induction. Suppose that $\neg\{E\}\phi$. A’s strategy then arises as follows. If E is to move, then no successor available to her can guarantee a win, since she has no winning strategy now - and so A can just “wait and see”. If A is to move, then there must be at least one possible move leading to a state where E has no winning strategy: otherwise, E has a winning strategy right now after all. Continuing this way, A is bound to produce runs of the kind described.²⁷

How would our earlier analysis extend to a setting like this, where infinite histories themselves are the outcomes of the game, and players try to achieve global properties of these?

Dynamics in games with imperfect information Moreover, many if not most games have *imperfect information*, with uncertainties for players where they are in the game tree. Think of card games, or other games with restricted observation. Can’t our analysis be extended to this area, where in general, Backward Induction no longer works? We merely illustrate the task ahead with two simple scenarios for the reader to ponder:

Example 9.2. (Strategic reasoning in imperfect information games)

In the following games, outcome values are written in the order (“A-value, E-value”):



The game to the left seems a straightforward extension of techniques for removing dominated moves, but that to the right raises tricky issues of what A would be telling E by moving right. We leave the question what should or will happen in both games to the reader: [21], [45] have more discussion.²⁹

10. Coda: alternatives to Backward Induction and true game dynamics

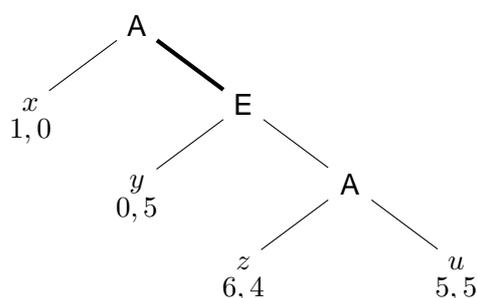
Finally, we return to our running example of Backward Induction. It has been claimed that its very style of reasoning is incoherent (cf. [15]):

The paradox of Backward Induction

Example 10.1. (The “Paradox of Backward Induction”)

Recall the style of reasoning toward a Backward Induction solution, as in:

²⁷This argument has a *co-algebraic* flavor, cf. [42], that we do not pursue here.
²⁸The tree to the right is adapted from an example in an invited lecture by Robert Stalnaker at the Gloriclass Farewell Event, ILLC Amsterdam, January 2010.
²⁹Further challenges to our analysis include equilibria with *coalitions* of players, and *simultaneous moves*. Nothing in our logics prevent this: it has just not been done yet.



Backward Induction tells us that A will go left at the start, on the basis of logical reasoning that is available to both players. But then, if A plays right (see the black line) what should E conclude? Does not this mean that A is not following the BI reasoning, and that all bets are off as to what he will do later on in the game? It seems that the very basis for the above computation collapses.

Responses to this conceptual difficulty vary - and many authors doubt that there is a genuine paradox here. The characterization result of [1] assumes that players know that rationality prevails throughout, something that [38] calls “rationality no matter what”, a stubborn unshakable belief that players will act rationally later on, even if they have not done so up until now.³⁰ [3] essentially take the same tack, deriving the BI strategy from an assumption of “stable true belief” in Rationality, a gentler form of stubbornness stated in terms of dynamic-epistemic logic.

Logics of actions, preference, and agent types Personally, we are more inclined toward another analysis, in line with [41]. A richer game analysis should add an account of the types of agent that play a game. In particular, we need to represent the *belief revision* policies by the players, that determine what they will do when making a surprising observation contradicting their beliefs in the course of a game. There are many different options for such policies in the above example, such as

- “It was just an error, and A will go back to being rational”,
- “A is inviting me to go right, and I will be rewarded for that”,
- “A is an automaton with a general rightward tendency”, and so on ...³¹

Our logical analysis so far omits this type of information about players of the game, since our algorithms make implicit uniform assumptions about what they are going to do as the game proceeds³².

Belief revision policies are not an explicit part of our models so far. Thus, our fixed-point logics tell only a limited story. Eventually, we may need a richer mathematical model for game solution, that can also deal with the dynamics of how players update knowledge and revise beliefs as a game unfolds.

³⁰One can defend this by assuming that the other player only makes isolated “mistakes”.

³¹One reaction to these surprise events might even be a switch to an entirely new reasoning style about the game. That might require more finely-grained *syntax-based* views of revision.

³²Some ideas in this direction occur in the area of computer game analysis, in the form of “opponent modeling”: [32].

11. Conclusion and general outlook

We have shown how standard logical fixed-point languages can define game solution procedures and their resulting relations of “best action”. We think that this is a good format for more general studies of game-theoretic notions, including finding alternatives to currently received views. But also, we hope to have shown that the game arena poses interesting problems for existing logics of computation, as one adds further structure that is typical for agents: preference, information, and eventually, even “processing types” for agents. All these contacts may eventually lead to legitimate children of logic and game theory.

References

- [1] Aumann, R.: Backward Induction and Common Knowledge of Rationality, *Games and Economic Behavior*, **8**, 1995, 6–19.
- [2] Baltag, A., Smets, S.: Group Belief Dynamics under Iterated Revision: Fixed-Points and Cycles of Joint Upgrades, *Proceedings of TARK 2009*, 2009, 41–50.
- [3] Baltag, A., Smets, S., Zvesper, J.: Keep “Hoping” for Rationality: A Solution to the Backward Induction Paradox, *Synthese*, **169**(2), 2009, 301–333.
- [4] van Benthem, J.: Extensive Games as Process Models, *Journal of Logic, Language and Information*, **11**, 2001, 289–313.
- [5] van Benthem, J.: Games in Dynamic Epistemic Logic, *Bulletin of Economic Research*, **53**(4), 2002, 219–248.
- [6] van Benthem, J.: Minimal Predicates, Fixed-Points, and Definability, *Journal of Symbolic Logic*, **70**(3), 2005, 696–712.
- [7] van Benthem, J.: Modal Frame Correspondences and Fixed-Points, *Studia Logica*, **83**(1), 2006, 133–155.
- [8] van Benthem, J.: Dynamic Logic of Belief Revision, *Journal of Applied Non-Classical Logics*, **17**(2), 2007, 129–155.
- [9] van Benthem, J.: Rational Dynamics and Epistemic Logic in Games, *International Game Theory Review (IGTR)*, **09**(01), 2007, 13–45.
- [10] van Benthem, J.: In Praise of Strategies, to appear in J. van Eijck and R. Verbrugge, eds., *Games, Actions and Social Software*, College Publications, London, 2008.
- [11] van Benthem, J.: The Logic of Rational Agency, Invited Lecture, Third Indian ICLA Winter School, Chennai, 2009.
- [12] van Benthem, J.: *Logical Dynamics of Information and Interaction*, Cambridge University Press, Cambridge, 2010.
- [13] van Benthem, J., van Otterloo, S., Roy, O.: *Modality Matters: Twenty-Five Essays in Honour of Krister Segerberg*, vol. 53 of *Philosophical Studies*, chapter Preference Logic, Conditionals, and Solution Concepts in Games, Lagerlund, Lindström and Sliwinski, Uppsala, 2006, 61–76.
- [14] van Benthem, J., Pacuit, E.: The Tree of Knowledge in Action, *Proceedings Advances in Modal Logic*, 2006, 87–106.
- [15] Bicchieri, C., Schulte, O.: Common Reasoning about Admissibility, *Erkenntnis*, **45**, 1997, 229–325.
- [16] Blackburn, P., de Rijke, M., Venema, Y.: *Modal Logic*, Cambridge University Press, Cambridge, 2000.

- [17] Bonanno, G.: The Logic of Rational Play in Games of Perfect Information, *Economics and Philosophy*, **7**, 1991, 37–65.
- [18] Boolos, G.: *The Logic of Provability*, Cambridge University Press, Cambridge, 1993.
- [19] de Bruin, B.: *Explaining Games: on the Logic of Game Theoretic Explanations*, Ph.D. Thesis, ILLC, University of Amsterdam, 2004.
- [20] Dawar, A., Grädel, E., Kreutzer, S.: Inflationary Fixed-Points in Modal Logic, *ACM Transactions on Computational Logic*, **5**, 2004, 282–315.
- [21] Dégrement, C.: *The Temporal Mind. Observations on the Logic of Belief Change in Interactive Systems*, Ph.D. Thesis, ILLC, University of Amsterdam, 2010.
- [22] Denecker, M., Ternovska, E.: A Logic of Nonmonotone Inductive Definitions, *ACM Trans. Comput. Logic*, **9**(2), 2008, 1–52.
- [23] van Ditmarsch, H., van der Hoek, W., Kooi, B.: *Dynamic-Epistemic Logic*, Cambridge University Press, Cambridge, 2007.
- [24] Ebbinghaus, H.-D., Flum, J.: *Finite Model Theory*, Springer, Berlin, 1995.
- [25] Geanakoplos, J., Polemarchakis, H.: We Can't Disagree Forever, *Journal of Economic Theory*, **28**, 1982, 192–200.
- [26] Gheerbrant, A.: *Fixed-Point Logics on Trees*, Ph.D. Thesis, ILLC, University of Amsterdam, 2010.
- [27] Gurevich, Y., Shelah, S.: Fixed-Point Extensions of First-Order Logic, *Annals of Pure and Applied Logic*, **32**, 1986, 265–280.
- [28] Halpern, J., Vardi, M.: The Complexity of Reasoning about Knowledge and Time, I: Lower Bounds, *Journal of Computer and System Sciences*, **38**(1), 1989, 195–237.
- [29] Harrenstein, P.: *Logic in Conflict*, Ph.D. Thesis, Institute of Computer Science, University of Utrecht, 2004.
- [30] van der Hoek, W., Pauly, M.: Modal Logic for Games and Information, in: *Handbook of Modal Logic* (P. Blackburn, F. Wolter, J. van Benthem, Eds.), Elsevier, Amsterdam, 2006, 1077–1148.
- [31] Immermann, N.: *Descriptive Complexity*, Springer, Berlin, 1999.
- [32] Iida, H., Uiterwijk, J.W.H.M., Herik, H.J. van den, and Herschberg, I.S. (1993). Potential Applications of Opponent-Model Search. Part 1: The Domain of Applicability. ICCA Journal, Vol. 16, No. 4, pp. 201-208.
- [33] Kreutzer, S.: Expressive Equivalence of Least and Inflationary Fixed-Point Logic, *Annals of Pure and Applied Logic*, **130**(1–3), 2004, 61–78.
- [34] Liu, F.: *Changing for the Better, Preference Dynamics and Agent Diversity*, Ph.D. Thesis, Ilc, University of Amsterdam, 2008.
- [35] Osborne, M., Rubinstein, A.: *A Course in Game Theory*, The MIT Press, Cambridge (Mass.), 1994.
- [36] van Otterloo, S.: *A Strategic Analysis of Multi-Agent Protocols*, Ph.D. Thesis, ILLC, Amsterdam, 2005.
- [37] Pacuit, E., Roy, O.: *Interactive Rationality*, Lectures Notes, University of Groningen and University of Tilburg, 2010.
- [38] Samet, D.: Counterfactuals in Wonderland, *Games and Economic Behavior*, **51**(2), 1997, 537–541.
- [39] Selten, R.: Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit, *Zeitschrift für die gesamte Staatswissenschaft*, **121**, 1965, 301–324.

- [40] Selten, R.: Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games, *International Journal of Game Theory*, **4**, 1975, 25–55.
- [41] Stalnaker, R.: Extensive and Strategic Form: Games and Models for Games, *research in Economics*, **53**, 1999, 293–319.
- [42] Venema, Y.: Algebras and Coalgebras, in: *Handbook of Modal Logic* (P. Blackburn, F. Wolter, J. van Benthem, Eds.), Elsevier, 2006, 331–426.
- [43] von Wright, G. H.: *The Logic of Preference*, Edinburgh University Press, Edinburgh, 1963.
- [44] Zvesper, J.: Arriving at the BI Condition, Manuscript, Oxford Computing Lab, Oxford University, 2010.
- [45] Zvesper, J.: *Playing with Information*, Ph.D. Thesis, ILLC, University of Amsterdam, 2010.