
8 Deliberation as Iterated Update

In Chapter 2 of this book, we took the Backward Induction procedure as our pilot example, and described it in a fixed point logic of action and preference, and also, zooming out to a less fine-grained level of detail, in a modal logic of best action. But there is also another way of construing the logical import of solving games. Backward Induction is a *procedure* for creating expectations about how a game will proceed. In the logical dynamics of the present part, it is the procedure itself that deserves attention, as a specimen of what we might call the dynamics of deliberation. Taking this view of game solution as rational procedure also has a broader virtue. We can do justice to an appealing intuition behind the fundamental notion of rationality: it is not a static state of grace, but a style of doing things. Along with this, we get a dynamic focus shift in the epistemic foundations of game theory that may be of interest in itself.

8.1 Backward Induction and announcing rationality

Let us shift attention from the static rationality in Part I to what players do when deliberating about a game. We will see how the relevant procedures reach stable limit models where rationality has become common knowledge. We show this for the method of Backward Induction, as a pilot for our style of analysis. The version of the algorithm we have in mind uses the relational strategies of Chapter 2.⁸⁵

⁸⁵ Indeed, the very word solution has an ambiguity between a reading as a procedure (“Solution of this problem is not easy”) and a static product of such a procedure (“Show me your solution”).

Our dynamic analysis of Backward Induction takes it to be a process of prior deliberation about a game by players whose minds proceed in harmony. The steps driving the information flow in our first scenario are public announcements $!\varphi$ saying that proposition φ is true, as explained in Chapter 7. These transform an epistemic model \mathbf{M} into its submodel $\mathbf{M}|\varphi$ whose domain consists of just those worlds in \mathbf{M} that satisfy φ .

The driver: Rationality We now explain the driving assertion about games that we will need in what follows. It is closely related to our analysis of best action in Chapter 2, though it differs slightly from the principle RAT used there.

DEFINITION 8.1 Node rationality

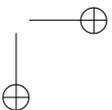
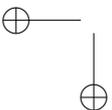
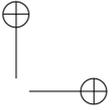
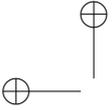
At a turn for player i in an extensive game, a move a is *dominated* by a sibling b (a move available at the same node) if every history through a ends worse, in terms of i 's preference, than every history through b . Now *rationality* (**rat**, for short) says that “at the current node, no player has chosen a strictly dominated move in the past coming here.” ■

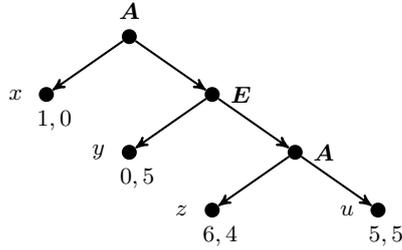
This makes an assertion about nodes in a game tree, namely, that they did not arise through playing a dominated move. This is often called playing a “best response.” Some nodes will satisfy this, others may not: we only need that **rat** is a reasonable local property of nodes. Thus, announcing this formula as a true fact about behavior of players is informative, and it will in general make a current game tree smaller.

But then we get a dynamics as in the earlier scenario of the Muddy Children in Chapter 7, where *repeated true assertions* of ignorance eventually produced enough information to solve the 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 always reaches a limit, a smallest subgame where no node is dominated any longer.

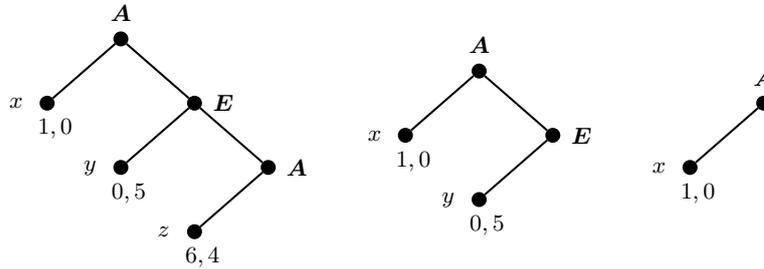
EXAMPLE 8.1 Solving games through iterated assertions of rationality

Consider the following extensive game with three turns, four branches, and payoffs for players **A** and **E** in that order:





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



Thus, the Backward Induction solution emerges step by step. ■

It is shown in van Benthem (2007d) that the actual Backward Induction path for extensive games is obtained by repeated announcement of the assertion *rat* to its limit. We will now explain this in more detail.

Logical background We reiterate some relevant notions from Chapter 7.

DEFINITION 8.2 Announcement limit

For each epistemic model M and each proposition φ true or false at points in M , the *announcement limit* $\#(M, \varphi)$ is the first model reached by successive announcements $!\varphi$ that no longer changes after the last announcement is made. That such a limit exists is clear for finite models, since the sequence of successive submodels is non-increasing, but announcement limits also exist in infinite models, where we stipulate that, at limit ordinals, intersections are taken of all previous stages. ■

There are two cases for the limit model. Either it is non-empty, and *rat* holds in all nodes, meaning it has become common knowledge (the “self-fulfilling” case),

or the limit model is empty, meaning that the negation $\neg\mathbf{rat}$ has become false (the “self-refuting” case). Both possibilities occur in concrete puzzles, although generally speaking, rationality assertions such as \mathbf{rat} tend to be self-fulfilling, while the ignorance statement that drives the Muddy Children was self-refuting: at the end, it held nowhere.

Capturing Backward Induction by iterated announcement With general relational strategies, the iterated announcement scenario produces the relational version of Backward Induction defined in Chapter 2.

THEOREM 8.1 In any game tree \mathbf{M} , $\#(\mathbf{M}, \varphi)$ is the actual subtree computed by Backward Induction.

Proof This can be proved directly, but it also follows from a few simple considerations. For a start, it turns out that it is easier to change the definition of the driving assertion \mathbf{rat} . 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 forests record more information.

Now we make some simple, but useful auxiliary observations.

Sets of nodes as relations Two simple facts about game trees are as follows.

FACT 8.1 Each subrelation R of the total *move* relation has a unique matching set of nodes $\mathit{reach}(R)$ being the set-theoretic range of R plus the root of the tree.⁸⁶

FACT 8.2 Vice versa, each set X of nodes has a unique corresponding subrelation of the move relation $\mathit{rel}(X)$ consisting of all moves in the tree that end in X .

These facts link the approximation stages BI^k for Backward Induction (i.e., the successive relations computed by our procedure in Chapter 3) and the stages of our public announcement procedure. They are in harmony all the way.

FACT 8.3 For each k , in each game model \mathbf{M} , $BI^k = \mathit{rel}((\mathbf{!rat})^k, \mathbf{M})$.

Proof The argument is by induction on k . The base case is obvious: \mathbf{M} is still the whole tree, and the relation BI^0 equals *move*. Next, consider the inductive step. If we announce \mathbf{rat} again, we remove all points reached by a move that is dominated

⁸⁶ Here, the root of the tree is only added for technical convenience.

for at least one player. Clearly, these are just the moves that were cancelled by the corresponding step of the Backward Induction algorithm. ■

It also follows that, for each stage k ,

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

Either way, we conclude that the algorithmic fixed point definition of the Backward Induction procedure and our iterated announcement procedure amount to the very same thing. ■

One might say that our deliberation scenario is just a way of conversationalizing the mathematical fixed point computation of Chapter 2. 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 model: it has made itself into *common knowledge*.

8.2 Another scenario: Beliefs and iterated plausibility upgrade

Next, in addition to knowledge, consider the equally fundamental notion of belief. Many foundational studies in game theory 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 today. We first recall the logical analysis of Backward Induction given in Chapter 2, and relate it to this perspective.

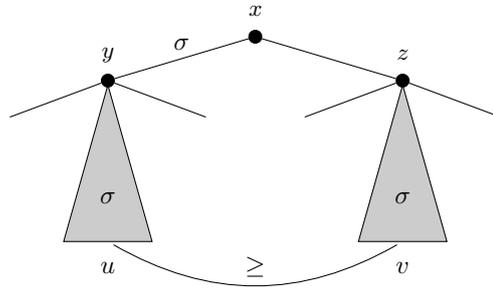
THEOREM 8.2 On finite extensive games, the *BI* strategy is the largest subrelation σ of the total *move* relation that has at least one successor at each node, while satisfying the following property for all players i :

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

This rationality assumption was a confluence property for action and preference:

$$CF \quad \forall x \forall y \left((Turn_i(x) \wedge x \sigma y) \rightarrow \forall z (x \text{ move } z \rightarrow \exists u \exists v (end(u) \wedge end(v) \wedge y \sigma^* u \wedge z \sigma^* v \wedge v \leq_i u)) \right)$$

that could be pictured in the following game tree with additional structure:



The shaded area is the part that can be reached via further play with our strategy. In the consequent of the syntactic $\forall\forall\exists\exists$ format for CF, all occurrences of σ are positive, and this was the basis for its definability in the standard first-order fixed point logic LFP(FO).

THEOREM 8.3 The *BI* relation is definable in LFP(FO).

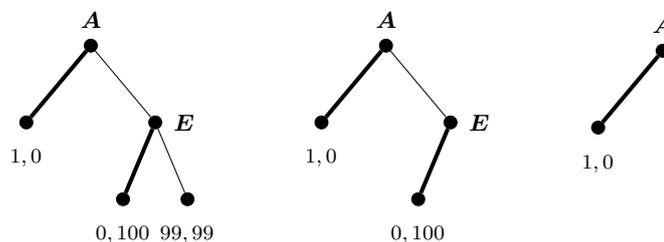
This connected game solution with fixed point logics of computation.

The important conceptual point here is that we need not think of the shaded parts as coming from further play by the strategy under consideration. We can also view them as the further histories that players think most plausible, encoding their expectations about the future. This gives us a connection with the dynamic logics of belief in Chapter 7.

Backward Induction in a soft light An appealing take on the relational Backward Induction strategy in terms of beliefs uses soft update that does not eliminate worlds like announcements $!\varphi$, but rearranges the plausibility order between worlds. A typical soft update in Chapter 7 was the radical upgrade $\uparrow\varphi$ that makes all current φ -worlds best, puts all $\neg\varphi$ -worlds underneath, while keeping the old ordering inside these two zones. Now recall our observation that Backward Induction creates expectations for players. The 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.

EXAMPLE 8.2 A debatable outcome, hard version

Consider the game that started Part I as a conceptual appetizer. The hard scenario in terms of events *!rat* removes nodes from the tree that are strictly dominated by siblings as long as this can be done, resulting in the following stages:



This scenario gives us players' absolute beliefs, but not yet conditional beliefs about what might have happened during off-path play. ■

By contrast, a soft scenario does not remove nodes but modifies the plausibility relation. We start with all endpoints of the game tree incomparable with respect to plausibility.⁸⁷ Next, at each stage, we compare sibling nodes, using an appropriate notion of rationality in beliefs.

DEFINITION 8.3 Rationality in beliefs

A move x for player i *dominates* its sibling y *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***, for short) is the assertion that no player plays a move that is dominated in beliefs. ■

Now we perform a relation change that is like an iterated radical upgrade ($\uparrow\varphi$).

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

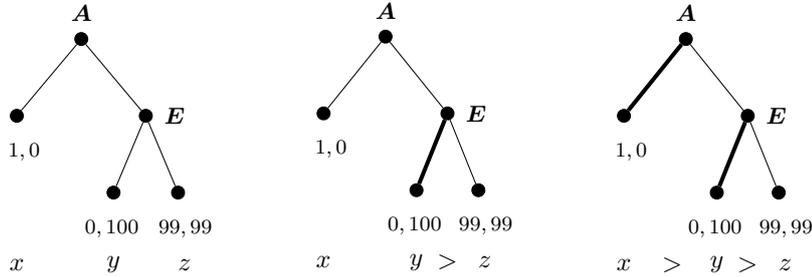
This changes the plausibility order, and hence the dominance pattern, so that an iteration can start.⁸⁸

EXAMPLE 8.3 A debatable outcome, soft version

The stages for the soft procedure in the above example are as follows, where we use the letters x , y , and z to stand for the end nodes or histories of the game:

⁸⁷ Other versions of our analysis would have all end nodes initially equiplausible.

⁸⁸ We omit some details; in general, the plausibility upgrades take place in subtrees only.



In the first tree, going right is not yet dominated in beliefs for A by going left. The assertion *rat** only has force at E 's turn, and update makes $(0, 100)$ more plausible than $(99, 99)$. After this change, however, going right has become dominated in beliefs, and a new update takes place, making A 's going left most plausible. ■

THEOREM 8.4 On finite trees, the Backward Induction strategy is encoded in the plausibility order for end nodes created in the limit 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 this result, using an idea from Baltag et al. (2009).

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

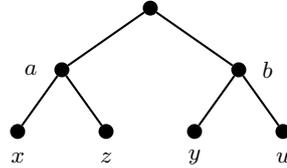
DEFINITION 8.4 Leaf order from a sub-move relation
 We put $x ord(R) y$ iff, looking upward at the first node z where the histories of x , y diverged, if x was reached via an R move from z , then so is y . ■

The following property of this order is easy to see by inspection of trees.

FACT 8.4 The relation $ord(R)$ is a total pre-order on leaves.

Moreover, this total order \leq on leaves is "tree-compatible," meaning that, for any two leaves x and 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 \leq to all leaves y' reached by taking the move toward y . This means that there can be no crisscrossing as in the following tree:





with $x < y < z < u$

DEFINITION 8.5 Relational strategies from leaf order

Conversely, any tree-compatible total order \leq on leaves induces a subrelation $rel(\leq)$ 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 maps rel and ord give a precise meaning to the way in which Baltag et al. (2009) can say that strategies are the same as plausibility relations.⁸⁹

Now we can relate the computation in our upgrade scenario for belief and plausibility to the earlier relational algorithm for Backward Induction. Things are in harmony at each stage.

FACT 8.5 For any game tree M and any k , $rel((\uparrow \mathbf{rat}^*)^k, M) = BI^k$.

Proof The key point was demonstrated in our example of a stepwise solution. When computing a next approximation for the Backward Induction relation according to CF, we drop those moves that are dominated in beliefs by another available one. This has the same effect as making those leaves that are reachable from dominated moves less plausible than those reachable from surviving moves. And that was precisely the earlier upgrade step. ■

Thus, the algorithmic analysis of Backward Induction and its procedural analysis in terms of forming beliefs amount to the same thing. Still, as with iterated announcements, the iterated upgrade scenario has interesting features of its own. One is that, for logicians, it accounts for the genesis of the plausibility orders usually treated as primitives in doxastic logic. Thus, games provide an underpinning for a possible world semantics of belief that is of independent interest.⁹⁰

89 Zvesper (2010) relates our dynamic analysis to achieving the sufficient condition for the Backward Induction outcome given in Baltag et al. (2009).

90 We have stated the operations ord and rel semantically. They can also be viewed as syntactic translations, and then various logical definitions for Backward Induction can be directly transformed into each other (see Gheerbrant 2010 for details).

8.3 Repercussions and extensions

Extensional equivalence, intensional differences Putting together the results in Chapter 2 and those found here, 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. Still, extensionally equivalent definitions can have interesting intensional differences in terms of what they suggest. For instance, the above analysis of strategy creation and plausibility change illustrates a general conceptual issue: the deep entanglement of agents’ beliefs and actions in the foundations of decision and game theory.

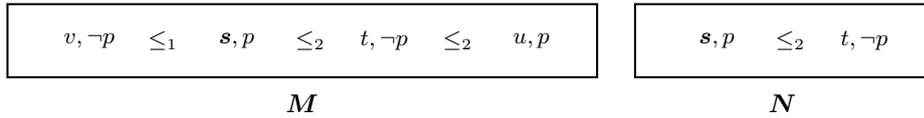
Dynamic instead of static foundations As we have said already, a key feature of our dynamic announcement and upgrade scenarios is that they are self-fulfilling: ending in non-empty largest submodels where players have common knowledge or common belief of rationality. Thus, this dynamic style of analysis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

Announcing rationality, hard or soft, is not the only case of interest. Deliberation can be driven by other statements, and it is not the only activity that falls under this way of thinking.

Other game-theoretic construals We have seen in Chapter 2 how Backward Induction can be viewed as producing various Nash equilibria between functional strategies with unique outputs at nodes. These sharper predictions of behavior corresponded to making assumptions about the player one is against: say, less or more careless about an opponent’s interests once the player’s own interests have been served. One obvious question is how to extend our current style of analysis to this setting. One way of doing this is by making such assumptions about players explicit, as we will do in Chapters 9 and 10.

Iterated hard announcements Our scenarios have a much broader sweep than may appear from our specific case study. Dégremon & Roy (2009) give a limit analysis for agents that communicate disagreement (cf. Aumann 1976) via iterated hard public announcements of conflicts in belief. They find interesting new scenarios, including the following one.

EXAMPLE 8.4 Stated disagreements in beliefs can switch truth value
 Consider two models M and N with actual world s , and accessibility running as indicated. For example, in M , $\{s, t, u\}$, $\{v\}$ are epistemic equivalence classes for agent 2 ordered by plausibility. For agent 1 , the epistemic classes are $\{s, v\}$, $\{t\}$, and $\{u\}$. Now, in M , $B_1\neg p \wedge B_2p$ is true at s and t only. Announcing it updates M to N , in whose actual world $B_1p \wedge B_2\neg p$ is true:

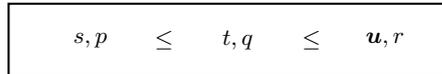


Here, truth values are computed entirely following our earlier clauses for belief. ■

This leads to a dynamic-epistemic version of the results in Geanakoplos & Polemarchakis (1982). Any dialogue where agents keep stating whether or not they believe that formula φ is true at the current stage leads to agreement in the limit. If agents share a well-founded plausibility order at the start (their hard information may differ), in the first fixed point, they all believe or all do not believe that φ is true. Dégrement (2010) links these results to syntactic definability of relevant assertions in epistemic fixed point logics.

Iterated soft updates Baltag & Smets (2009) analyze limit behavior of soft announcements, including the radical $\uparrow \varphi$ and conservative $\uparrow \varphi$ of Chapter 7. Surprises occur, and their flavor is given in the following illustration.

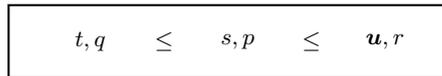
EXAMPLE 8.5 Cycling radical upgrades
 Consider a one-agent plausibility model with proposition letters as indicated:



Here u is the actual world. Now make the following soft announcement:

$$\uparrow (r \vee (B^{-r}q \wedge p) \vee (B^{-r}p \wedge q))$$

The formula is true in worlds s and u only, and hence the new pattern becomes:



In this new model, the formula $r \vee (B^{-r}q \wedge p) \vee (B^{-r}p \wedge q)$ is true in t and u only, so radical upgrade returns the original model, starting a cycle:

$$s, p \leq t, q \leq u, r$$

Note that the final world u always stays in place in this scenario. ■

This example provides a formal modeling for two important phenomena. Oscillation in public opinion is a fact of social information flow, which definitely needs further logical study. But the stability of the final world is significant, too. Baltag & Smets (2009) prove that, despite cycles with conditional beliefs, every truthful iterated sequence of radical upgrades stabilizes all absolute factual beliefs. This stabilization result is relevant to the next application of the above ideas.

From belief revision to learning theory Iterated plausibility upgrade also applies to “learning in the limit” as studied in formal learning theory (Kelly 1996), leading to interesting results in Baltag et al. (2011) that throw more light on the earlier mechanism. The set of hypotheses at stake in a learning problem creates an initial epistemic model over which a set of all possible, finite or infinite, histories of signals is then given to the learner, as procedural information about the process of inquiry. The aim of the learning is to find out where the actual history lies in some given partition, corresponding to an issue as discussed in Section 7.7 of Chapter 7. Observing successive signals then triggers either hard or soft information about the actual history.

A key observation is that learning methods can be encoded as plausibility orderings on the initial model that determine agents’ beliefs about the issue in response to new input. It then turns out that both iterated public announcement and iterated radical upgrade are universal learning methods, although only radical upgrade still has this feature in the presence of (finitely many) errors in the input stream.⁹¹

⁹¹ Gierasimczuk (2010) has details, including definability results for finite identifiability and identifiability in the limit in the epistemic-temporal language of Chapter 5. One natural question still open is to what extent the initial plausibility order can be mimicked by belief formation on the fly in the process of receiving signals.

8.4 Logical aspects

The preceding topics raise a number of general logical questions. While these may not be relevant to specific games as such, they do provide a broader setting for the kind of analysis that we have proposed in this chapter.

Fixed point logics While relational Backward Induction was definable in the first-order fixed point logic, LFP(FO), this depended on positive occurrence of σ in the syntax of the assertion CF, as noted above. To test the scope of our method, consider a natural maximin variant BI[#] of the Backward Induction algorithm where choices between moves ensure the greatest minimal value. We leave it to the reader to see how the following rule can deviate from our more cautious relational Backward Induction algorithm. This time, the syntactic confluence property CF[#] is

$$\bigwedge_i (Turn_i(x) \rightarrow \forall y(x\sigma y \rightarrow (x \text{ move } y \wedge \forall u((end(u) \wedge y\sigma^*u) \rightarrow \forall z(x \text{ move } z \rightarrow \exists z(end(v) \wedge zS\sigma^*v \wedge v \leq_i u))))))$$

where not all occurrences of the relation symbol S are positive. Hence, CF[#] cannot be used for an immediate fixed point definition in LFP(FO). But we do have a characterization in a slightly extended logical system.

THEOREM 8.5 The relational BI[#] strategy is definable in “first-order deflationary fixed point logic” IFP(FO) using simultaneous fixed points.

Proof A proof can be found in van Benthem & Gheerbrant (2010). ■

Unlike the systems discussed in Chapters 1 and 2, deflationary fixed point logic puts no restrictions on the formulas $\varphi(P)$ used in fixed point operators, but it forces convergence from above by always intersecting the new set with the current approximation. This system is of major interest in understanding computation, but we refer the reader for details to Ebbinghaus & Flum (1999) and Dawar et al. (2004).⁹² We will discuss this logic further in Chapter 13 when analyzing solution procedures for strategic games.

⁹² By the results of Gurevich & Shelah (1986) and Kreutzer (2004), BI[#] is still definable in LFP(FO) by using extra predicates. However, their computation no longer matches stages of our algorithm.

Exploiting the well-foundedness of trees Fixed point logics such as LFP(FO) or IFP(FO) work on any model. This generality is attractive when investigating abstract solution procedures for classes of games. However, another approach is possible. Variants of Backward Induction exploit a special feature of finite extensive games, namely, their *well-founded* tree dominance order.⁹³ Such orders allow recursive definitions without positive occurrence as long as all occurrences of the defined predicate scope under quantifiers looking downward along the ordering.⁹⁴ Thus, we get many more recursive definitions on game trees (see Gheerbrant 2010 for matching logics of trees).

Finally, fixed point logics like the above are also conceptually intriguing from the perspective of statics versus dynamics raised in the transition from Part I to Part II. They have a bit of both, since their fixed point operators come with procedures.

Limits in dynamic-epistemic logic We have already seen how limit scenarios for game solution and related tasks such as conversation or learning raise interesting logical issues. We noted in Chapter 7, and again in this chapter, that iterated announcements can end in limit models $\#(\mathbf{M}, \varphi)$ where for the first time, a new event $!\varphi$ no longer changes things. These models came in two kinds, non-empty $\#(\mathbf{M}, \varphi)$, where φ has become common knowledge, and models where φ has become false in the actual world. Rationality assertions *rat* were of the former self-fulfilling kind, while the ignorance statement driving the Muddy Children was of the latter self-refuting kind. Likewise, announcements of disagreement were self-refuting in Dégrement & Roy (2009). Can we say more? Going beyond the few known examples in games and elsewhere, can we say something systematic about the outcome from the syntactic form of the statements and the shape of the initial model?

A simpler related issue is the Learning Problem for public announcement (van Benthem 2011d). Using the notions in Chapter 7, it is easy to show that factual formulas become known upon announcement. But epistemic formulas need not behave in this manner. Moore sentences $p \wedge \neg Kp$ became false when announced truly. Thus, the problem arises of which syntactic shapes of dynamic-epistemic formulas φ guarantee that public announcement makes them known – i.e., $[\!|\varphi]K\varphi$ is valid. Holliday & Icard (2010) solve this problem.

93 Also, all trees allow for recursion over their predecessor ordering toward the root.

94 More precisely, $CF^\#$ defines its unique subrelation of the *move* relation by recursion on the well-founded tree order given by the relational composition of the sibling and dominance orderings.

It is an open problem to characterize the self-fulfilling and self-refuting dynamic-epistemic formulas φ syntactically. In fact, this behavior may be so dependent on the initial model, that uniform behavior is rare. Still, van Benthem (2007d) shows how limits of iterated public announcement on epistemic models are definable in a deflationary fixed point extension of the modal μ -calculus. Moreover, behavior gets better for “positive-existential formulas” constructed using this syntax:

$$\text{literals } (\neg)p \mid \wedge \mid \vee \mid \text{existential modalities } \diamond$$

FACT 8.6 Limit models for positive-existential modal formulas φ have their domain definable by a formula in the modal μ -calculus.

Proof The reason for the fixed-point definability is that positive-existential formulas have monotonic approximation maps in their announcement sequence. This will be covered in more detail in Chapter 13. ■

Both the rationality statement in our Backward Induction analysis and the ignorance statement in the Muddy Children problem are positive-existential. Fact 8.6 then shows that their logic remains simple and decidable. However, the disagreement statement of Dégrémont & Roy (2009) is not positive-existential, and yet its limit logic seems simple. We still do not understand in general why rationality is self-fulfilling, and disagreement self-refuting, on the above models.

All of these issues of limit definability and predicting behavior from syntax return with iterated upgrade of plausibility orderings. No general results seem to be known at this interface of dynamic logics and dynamical systems.

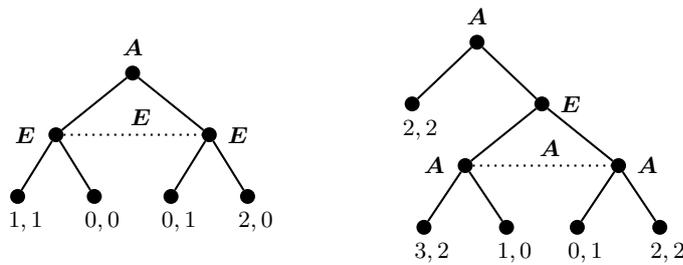
Fragments and complexity Moving from definability to proof, which logics are suited for reasoning with our dynamic scenarios? One relevant system seems public announcement logic with Kleene iteration added, PAL^* , but this system is highly complex. As noted in Chapter 7, Miller & Moss (2005) prove that validity in PAL^* is Π_1^1 -complete.

In addition to this source of high complexity, we saw in Chapter 2 that combinations of action and preference satisfying gridlike confluence properties can generate complexity as well. One way out here is that game solution procedures need not use the full power of logical languages for recursive procedures. Which fragments are needed? Moreover, PAL^* might be too ambitious, since we may just want to reason about limit models, not all intermediate stages. A second way out, mentioned in Chapter 7, is switching to more general temporal protocol models for dynamic-epistemic logics where complexity may drop.

At the moment, we are not sure what best dynamic logic to use for the theory of game solution, or more generally, a theory of protocols in temporal universes of informational events. The epistemic-temporal and dynamic logics in Fagin et al. (1995), van Benthem et al. (2009a), and Wang (2010) seem relevant, and provide lower-complexity tools for a wide array of tasks.⁹⁵

Infinite models Do our deliberation scenarios extend to infinite games? Infinite ordinal sequences are easy to add to iterations, and fixed point definitions make sense in infinite models. As we saw in Chapter 5, there may be game-theoretic substance to this generalization, since in infinite trees, intuitive reasoning changes direction from backward to forward. An illustration was our recursive analysis of weak determinacy. In this step, the mathematical spirit changed from inductive to co-inductive (Venema 2006), something that also proved attractive for strategies in Chapter 4.

Dynamics in games with imperfect information Many games have imperfect information, with uncertainties for players where they are in the tree. Can our dynamic analysis be extended to this area, where Backward Induction no longer works? We repeat an example from Chapter 3 that the reader may want to try. In the following games, outcome values are written in the order (*A*-value, *E*-value):



The game to the left yields to our technique of removing dominated moves, but the one to the right raises tricky issues of what *A* is telling *E* by moving right. Some of these issues will return in Chapter 9 on dynamics in models for games.

⁹⁵ Incidentally, while we have concentrated on modal formalisms here, all questions raised in this chapter make sense for other logical languages as well.

8.5 Conclusion

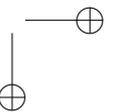
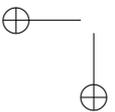
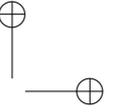
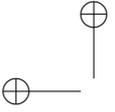
We have shown how the deliberation phase of games can be analyzed in terms of dynamic-epistemic iteration scenarios of knowledge update and belief revision. This style of thinking applies more broadly to epistemic puzzles, conversational scenarios, and learning methods. Beyond this conceptual contribution, our analysis raised new technical issues. We extended standard epistemic and doxastic logic with a notion of limit models, an intriguing topic that has hardly been explored yet. Our analysis also creates new bridges between game theory and fixed point logics, a natural mathematics of recursion that fits very well with equilibrium notions.⁹⁶

8.6 Literature

This chapter follows van Benthem (2007d) and van Benthem & Gheerbrant (2010).

Of related work going into more depth, we mention Baltag et al. (2009), Dégremont & Roy (2009), de Bruin (2010), Gheerbrant (2010), Baltag et al. (2011), and Pacuit & Roy (2011).

⁹⁶ We think that our scenarios will also provide a good format for developing alternatives to received views in game theory, although, admittedly, we have not done so here. Chapter 13 will give a few more examples, but even there we do not stray far from orthodoxy.

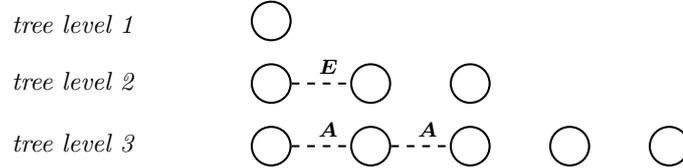


9 Dynamic-Epistemic Mechanisms at Play

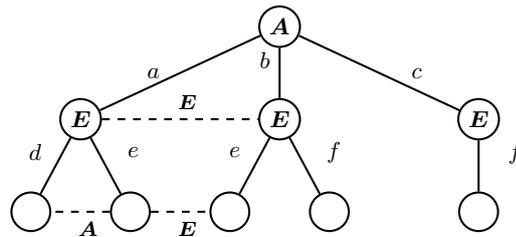
Games involve different sorts of dynamic events. The prior phase of deliberation was studied in Chapter 8 with the help of update mechanisms from Chapter 7 that change beliefs and create expectations. In this chapter, we turn to what happens during actual play. We will use the dynamic techniques of Chapter 7 once more, this time, to look at various kinds of events and information flow as a game proceeds. We do this first by making sense of a given record of a game, in particular, the uncertainty annotations found in imperfect information games. We will make their origins explicit in terms of dynamic-epistemic scenarios that produce these traces, first for knowledge, then also for belief. Next, we discuss updates during play, as the current stage keeps shifting forward. Our vehicle here is the epistemic-temporal forest perspective of Chapter 6 that encodes knowledge and belief for players of a game. While we viewed these models before as complete records of play, they can also serve as information states that can be modified by further events. Next, we show how our techniques can also analyze activities after play, as players ponder a game that has already happened, perhaps rationalizing what they did post facto. We also add some observations on more drastic events such as game change, since the same methods apply. In the course of this analysis, several issues will come to light about how all of these activities and events can work together harmoniously, although we will not present one unified theory. Some further thoughts on the overall program will be found in our concluding Chapter 10 in Part II.

9.1 Retrieving play from a game record

We start with the issue of making the dynamics explicit that lies behind a given game. As we saw in Part I, games annotated with imperfect information links are



The resulting annotated tree is the following imperfect information game:



This is a special case of “update evolution,” a process that creates successive epistemic models from an initial model by iterated update with some event model, or a sequence of event models. We will soon see in more detail how this works. ■

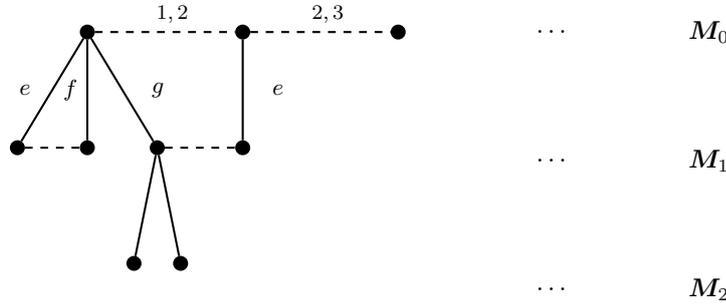
Game trees that are decorated with uncertainty links in this way are not arbitrary. There are special patterns, if the process worked in this systematic manner. We will now analyze what these are. In order to achieve a higher level of generality, we move from games to the more general epistemic-temporal models of Chapter 5. These models recorded information that agents have about the current protocol.

9.2 A representation for update on epistemic-temporal models

In update evolution, an initial epistemic model M is given, and it then gets transformed by the gradual application of event models $\mathcal{E}_1, \mathcal{E}_2, \dots$, to form a growing sequence of stages for an epistemic forest model

$$M_0 = M, \quad M_1 = M_0 \times \mathcal{E}_1, \quad M_2 = M_1 \times \mathcal{E}_2, \dots$$

It helps to visualize this in trees, or rather forest models such as the following:



Stages are horizontal, and worlds may extend downward via 0 or more successors. Through product update, worlds in the successive models arise from pair formation, resulting in sequences starting with one world in the initial model \mathbf{M} followed by a finite sequence of events that were executable when their turn came. Such worlds are essentially histories in the sense of the epistemic forest models of Chapter 5.

DEFINITION 9.1 Induced epistemic forests

Given a model \mathbf{M} and a finite or countable sequence of event models \mathbb{E} , the *induced epistemic forest model* $Forest(\mathbf{M}, \mathbb{E})$ has as its histories all finite sequences (w, e_1, \dots, e_k) produced by product update with successive members of the sequence \mathbb{E} , with accessibility relations and a valuation defined as in Chapter 7. ■

NOTE We will refer to epistemic forest models as ETL models henceforth, for the sake of brevity. We will also use the abbreviation DEL to remind the reader of dynamic-epistemic product update in the presence of partial observation for different agents.

Induced ETL models have a simple protocol \mathbb{H} (in the sense of Chapter 5) of available histories that determine how the total informational process can evolve, namely, only along the finite sequences that pass the requirements of the DEL update rule. The following three striking properties make these models stand out.

FACT 9.1 ETL models \mathbf{H} of the form $Forest(\mathbf{M}, \mathbb{E})$ satisfy the following three principles, where quantified variables h, h', k, \dots , range only over histories present in the initial model \mathbf{M} :

Perfect Recall If $he \sim k$, then there is some f with $k = h'f$ and $h \sim h'$.

Uniform No Miracles If $h \sim k$, and $h'e \sim k'f$, then $he \sim kf$.

Definable Execution The domain of any event e , viewed as a set of nodes in the forest model \mathbf{H} is definable in the epistemic base language.

The crucial observation is that these three properties induce the following representation theorem for dynamic-epistemic update evolution in epistemic-temporal forest models.⁹⁸

THEOREM 9.1 For ETL models \mathbf{H} , the following two conditions are equivalent:

- (a) \mathbf{H} is isomorphic to some model $Forest(\mathbf{M}, \mathbb{E})$.
- (b) \mathbf{H} satisfies Perfect Recall, Uniform No Miracles, and Definable Executability.

Proof The direction from (a) to (b) is given by Fact 9.1. Conversely, consider any ETL model \mathbf{H} satisfying the three conditions. Define an update sequence as follows:

- (a) The initial model \mathbf{M} consists of the set of histories in \mathbf{H} of length 1, copying their given epistemic accessibilities and valuation.
- (b) The event model \mathcal{E}_k is the set of events occurring at tree level $k + 1$ in \mathbf{H} , setting $e \sim f$ if there exist histories s and t of length k with $se \sim tf$ in \mathbf{H} . The required definability of event preconditions comes from Definable Executability.

We prove by induction that the tree levels \mathbf{H}_k at depth k of the ETL model \mathbf{H} are isomorphic to the successive epistemic models $\mathbf{M}_k = \mathbf{M} \times \mathcal{E}_1 \times \dots \times \mathcal{E}_{k-1}$.

The crucial fact is this, using our definition and the first two properties (writing (s, e) for the history se)

$$(s, e) \sim_{\mathbf{H}_k} (t, f) \quad \text{iff} \quad (s, e) \sim_{\mathbf{M}_k} (t, f)$$

We first proceed from left to right. By Perfect Recall, $s \sim t$ in \mathbf{H}_{k-1} , and therefore, by the inductive hypothesis, $s \sim t$ in \mathbf{M}_{k-1} . Next, by our definition of accessibility, $e \sim f$ in \mathcal{E}_k . Then, by the forward half of the DEL product update rule, it follows that $(s, e) \sim_{\mathbf{M}_k} (t, f)$.

Next, we proceed from right to left. By the other half of the definition of product update, $s \sim t$ in \mathbf{M}_{k-1} , and by the inductive hypothesis, $s \sim t$ in \mathbf{H}_{k-1} . Next, since $e \sim f$, by our definition, there are histories i and j with $ie \sim jf$ in \mathbf{H}_k . By Uniform No Miracles then, $se \sim tf$ holds in \mathbf{H} . ■

⁹⁸ Successive versions of this result have appeared in van Benthem (2001b), van Benthem & Liu (1994), and van Benthem et al. (2009a).

This result assumes linguistic definability for preconditions of events e , i.e., the domains of the matching partial functions in the tree \mathbf{H} .

There is also a purely structural version in terms of a notion from Chapter 1.

THEOREM 9.2 Theorem 9.1 still holds when we replace Definable Executability by Bisimulation Invariance: that is, closure of event domains under all purely epistemic bisimulations of the ETL model \mathbf{H} .

Proof Two facts from Chapter 1 suffice: (a) epistemically definable sets of worlds are invariant for epistemic bisimulations, and (b) each bisimulation-invariant set has an explicit definition in the infinitary version of the epistemic language.⁹⁹ ■

Our results state the essence of DEL update as a mechanism creating epistemic-temporal models. It is about agents with perfect memory, driven by observation only, whose information protocols involve only local epistemic conditions on executability of actions or events.

Caveat Our treatment implies “synchronicity”: uncertainty only occurs between worlds at the same tree level. Dégrement et al. (2011) present an important amendment showing how synchronicity is an artifact of the above representation that can be circumvented while keeping the spirit of the other principles, thereby allowing for DEL-induced epistemic forest models in which processes occur asynchronously.

Extended preconditions A mild relaxation of the above definability or invariance requirements for events allows preconditions that refer to the epistemic past beyond local truth. Think of a conversation that forbids repeated assertions: this protocol needs a memory of what was said, which need not be encoded in a local state.

Variety of players In Chapter 3, we looked at different kinds of agents, on a spectrum from perfect recall to memory-free. Our representation can be modified to characterize effects of other update rules, say, for agents with bounded memory.

EXAMPLE 9.2 Update for memory-free agents
A modified product update rule for completely memory-free agents works as follows:

$$(s, e) \sim (t, f) \text{ iff } e \sim f$$

Note how the prior worlds play no role at all, only the last event counts. ■

⁹⁹ This only guarantees finite epistemic definitions for preconditions in special models. However, further tightening of conditions has no added value in grasping the essentials.

Alternately, one can think not of agents that are memory-impaired, but of agents following a strategy that uses no memory. This is a general reinterpretation for our results, leaving the nature of the agents open, but capturing their styles of behavior.

9.3 Tracking beliefs over time

The preceding epistemic temporal analysis generalizes to other attitudes that are fundamental to rational agency, and especially, to beliefs that players have, based on their observations. The relevant structures are *epistemic-doxastic-temporal models* (DETL models), that is, branching forests as before, but with nodes in the same epistemic equivalence classes now also ordered by plausibility relations for agents. These expanded forest models interpret belief modalities at finite histories in the manner of Chapter 7. But as before, their belief relations can be very general, and as with knowledge, it makes sense to ask which of them arise as traces of some systematic update scenario.

REMARK Beliefs versus expectations

A clarification may be needed here. As we have seen in Chapter 6, intuitively, beliefs in a game come in two kinds: procedural beliefs about the game and its players, and expectations about the future. The scenario in this section is mainly about the first kind, that is, about beliefs where we stand. We will discuss connections with expectations, as created by Backward Induction or other deliberation methods, later on.

Following van Benthem & Dégrémont (2008), we take epistemic-doxastic models \mathcal{M} and plausibility event models \mathcal{E} to create products $\mathcal{M} \times \mathcal{E}$ whose plausibility relation obeys a notion from Baltag & Smets (2008) introduced in Chapter 7:

$$\text{Priority Rule} \quad (s, e) \leq (t, f) \text{ iff } (s \leq t \wedge e \leq f) \vee e < f$$

Let update evolution take place from some initial model along a sequence of plausibility event models $\mathcal{E}_1, \mathcal{E}_2, \dots$ according to some uniform protocol.¹⁰⁰ The crucial pattern that arises in the forest model created by the successive updates can be described as follows.

¹⁰⁰ Dégrémont (2010) also analyzes pre-orders, and “state-dependent” protocols where the sequence of event models differs across worlds of the initial epistemic model.

FACT 9.2 The histories h, h' and j, j' arising from iterated priority update satisfy the following two principles for any events e and f :

Plausibility Revelation Whenever $je \leq j'f$, then $he \geq h'f$ implies $h \geq h'$.

Plausibility Propagation Whenever $je \leq j'f$, then $h \leq h'$ implies $he \leq h'f$.

Together, these properties express the revision policy in the Priority Rule: its bias toward the last-observed event, but also its conservativity with respect to previous worlds whenever possible given the former priority.

THEOREM 9.3 A DETL model is isomorphic to the update evolution of an epistemic-doxastic model under successive epistemic-plausibility updates iff it satisfies the structural conditions of Section 9.2, with Bisimulation Invariance now for epistemic-doxastic bisimulations, plus Plausibility Revelation and Propagation.

Proof The idea of the proof is as before. Given a DETL forest \mathbf{H} , we say that

$e \leq f$ in the epistemic plausibility model \mathcal{E}_k if e, f occur at the same tree level k , and there are histories h and h' with $he \leq_{\mathbf{H}} h'f$.

One checks inductively, using priority update plus Plausibility Revelation and Plausibility Propagation in the forest \mathbf{H} , that the given plausibility order in \mathbf{H} matches the one computed by sequences of events in the update evolution stages

$$\mathbf{M}_{\mathbf{H}} \times \mathcal{E}_1 \times \dots \times \mathcal{E}_k$$

starting from the epistemic plausibility model $\mathbf{M}_{\mathbf{H}}$ at the bottom of the tree. ■

One can think of the structures described here as generalized imperfect information games, where information sets now also carry plausibility orderings.

Logical languages Languages over these models extend dynamic doxastic logic to a temporal setting as in Chapter 6. In particular, the safe belief modality of Chapter 7 is used in van Benthem & Dégrémont (2008) to state correspondences of Plausibility Revelation and Plausibility Propagation with special properties of agents. Dégrémont (2010) proves completeness for the logics, relating them to those used in Bonanno (2007), while also connecting doxastic protocols with formal learning theory. We will look at these logics in a different perspective in the next section.

9.4 Witnessing events and adding temporal logic

Having shown how dynamic actions may be retrieved from their traces in a game as usually defined, let us now place the focus directly on actual events that can take place during play.

The simplest example of such an event is the mere playing of a move, and its bare observation. In Chapter 10, we will take up the more sophisticated events discussed at the end of Chapter 6, interpreting moves as intentional, accidental, or otherwise. At this point, we ask how bare observation is dealt with in our dynamic logics. We expect an analysis close to our natural reading of the temporal models of Chapter 5, and this is borne out.

Technically, the topic to come is simpler than the scenarios in the preceding sections, since we focus on perfect information games with events that are publicly observed. Once this is understood, an extension to imperfect information games should be easy to make. Our setting is branching temporal models or epistemic forests, but for many points, the precise choice of models is immaterial. Moreover, we add a theme that was absent from Sections 9.2 and 9.3, namely, axiomatization in logical languages.

Playing a move involves change, as the current point of the forest model shifts, and this can be defined as follows.

DEFINITION 9.2 Updates for moves

An *occurrence* $!e$ of an event e changes the current pointed model (\mathbf{M}, s) to the pointed model (\mathbf{M}, se) , where the distinguished history s moves to se .¹⁰¹ ■

An equivalence with a standard modality is an immediate consequence.

FACT 9.3 The dynamic modality $\langle !e \rangle \varphi$ is equivalent to $\langle e \rangle \varphi$.

REMARK Existential modalities

Existential modalities often make logical principles in this area a bit easier to state. We will use them for this reason here and later in this chapter. Of course, in contexts with unique events, the difference between existential and universal modalities will be slight. Also, we use \diamond for the existential knowledge modality, suppressing agent

¹⁰¹ Here we assume, as is often done, that moves are fine-grained enough to be unique. Another view of this factual change is that the current history h gets extended to he .

indices for convenience. In the temporal language of Chapter 5 for forest models, the preceding event modality would be $F_e\varphi$, and this is in fact the notation that we will use in this section.

Moves under public observation are a very special case of the potentially much more intricate scenarios provided by DEL product update in preceding sections. Even so, it is illuminating to connect such bare events to earlier topics in this book. First, we show how principles of temporal logic reflect semantic properties of play, as counterparts to earlier laws of dynamic-epistemic logic.

FACT 9.4 The following principle is valid for knowledge and action on forest models with public observation of events:

$$F_e\Diamond\varphi \leftrightarrow (F_e\top \wedge \Diamond F_e\varphi)$$

We have shown the validity of this equivalence in Section 6.5 of Chapter 6. One can view this as a temporal equivalent of the PAL recursion axiom for knowledge, thinking of $!e$ as a public announcement that e has occurred. The precondition for this event is $F_e\top$, which fits with the protocol version of PAL in Chapter 7.

Representation once more The stated principle is not generally valid on all epistemic forest models. In the spirit of Section 9.2, it corresponds to the conjunction of two earlier properties:

Perfect Recall $\quad \forall xyz : ((xR_e y \wedge y \sim_i z) \rightarrow \exists u(x \sim_i u \wedge uR_e z))$
 (uncertainty after a move e can only come from earlier uncertainty).

No Miracles $\quad \forall xyz : ((x \sim_i y \wedge xR_e z \wedge yR_e u) \rightarrow z \sim_i u)$
 (uncertainty before a move e must persist after that same move,
 i.e., epistemic links can only be broken by different observations).

In slightly modified forms, these properties were also prominent in Chapters 3 and 5. For instance, agents with perfect recall will always know their past history. In their current form, they support a special case of the earlier representation theorem for epistemic forests, where the event models consist of isolated points. We leave the simple details of this specialization to the reader.

Other logical laws Interestingly, for general modal reasons, when events are unique as in this case, further laws will break down postconditions after events:

$$F_e(\varphi \wedge \psi) \leftrightarrow (F_e\varphi \wedge F_e\psi) \quad F_e\neg\varphi \leftrightarrow (F_e\top \wedge \neg F_e\varphi)$$

General product update What if events happen in a game according to the more general event models of Section 9.2? In that case, the logical axiom is the following, where we assume for convenience that only one event model \mathcal{E} was applied repeatedly, the way things worked in the imperfect information game of Section 9.1. What we get is essentially the characteristic DEL recursion axiom from Chapter 7.

FACT 9.5 On forest models produced by product update, the following is valid:

$$F_e \diamond \varphi \leftrightarrow (F_e \top \wedge \bigvee \{ \diamond F_f \varphi \mid e \sim f \text{ in } \mathcal{E} \})$$

Beliefs The same analysis applies to beliefs modeled by plausibility relations in DETL forests. We get temporal counterparts to earlier principles of dynamic plausibility change. First, we state some laws governing public events !e. These involve absolute and conditional belief, now in existential forms $\langle B \rangle \varphi$, $\langle B \rangle^\psi \varphi$, plus a modality $\langle \leq \rangle \varphi$ for safe belief, that can define the other two (cf. Chapter 7).

FACT 9.6 The following principles are valid on doxastic forest models:

- (a) $F_e \langle B \rangle^\psi \varphi \leftrightarrow (F_e \top \wedge \langle B \rangle^{F_e \psi} F_e \varphi)$
- (b) $F_e \langle \leq \rangle \varphi \leftrightarrow (F_e \top \wedge \langle \leq \rangle F_e \varphi)$

Proof We proved (a) in Section 6.9 of Chapter 6 in a model with plausibility running between histories, and (b) is even simpler in that setting. In forest models with plausibility relations between nodes, the argument is similar, using the definition of priority update in the special case when event models just have isolated points. ■

Once more, this result shows an earlier phenomenon: the technical similarity between history- and stage-based models. This shows in the following version for doxastic forests created by priority update with general event models \mathcal{E} . (The existential modality \diamond in our formula is epistemic over all \sim -accessible worlds.)

FACT 9.7 The following is valid on forest models created by priority update:

$$F_e \langle \leq \rangle \varphi \leftrightarrow (F_e \top \wedge \bigvee (\{ \langle \leq \rangle F_f \varphi \mid e \leq f \text{ in } \mathcal{E} \} \vee \{ \diamond F_f \varphi \mid e < f \text{ in } \mathcal{E} \}))$$

Note the analogy with the key recursion axiom for belief revision in Chapter 7.

This logic captures the preference propagation and preference revelation that characterized forest models of this sort. For instance, propagation said that, if $j e \leq j' f$, then $h \leq h'$ implies $h e \leq h' f$. This is expressed by the following temporal

formula with an existential modality E over the whole forest, and past modalities:

$$EF_e \langle \leq \rangle P_f^U \top \rightarrow (\langle \leq \rangle F_f \varphi \rightarrow [e] \langle \leq \rangle \varphi)$$

This can be derived from the preceding recursion principle as a law of the system.

9.5 Help is on the way: Hard information during play

Having reviewed simple events corresponding to official moves of a game, let us now consider a more ambitious scenario. Public moves are not the only events that occur during play. One may also experience events where further information comes in about the structure of the game, or the behavior of other players. There are many such scenarios, and we will discuss a few later on. This means that we now leave the official conception of a game, and we will reflect on this as we proceed.

The simplest new events are public announcements $!\varphi$ of information relevant to play. Here we will apply the standard view of world elimination from (\mathbf{M}, s) to $(\mathbf{M}|\varphi, s)$ to pointed forest models \mathbf{M} with finite histories as worlds. Using the methods of Chapter 7, we can then analyze a wide range of effects on earlier notions of action, knowledge, and belief. We note beforehand that this raises some delicate issues of interpretation, since forest models now become modifiable stages in a dynamic process, rather than universal receptacles of everything that might happen.

Recursion axioms for announcements in forests Effects of basic informational actions can be described explicitly on top of our static game languages.

THEOREM 9.4 The logic of public announcement in forest models is axiomatizable.

Proof As in Chapter 7, the heart of the analysis is finding the right recursion laws for the announcement modality $\langle !\varphi \rangle \psi$. We consider the various postconditions ψ that can occur. The recursion axioms for atoms and Boolean operators are as usual.

Action Consider the pure event structure of forest models (cf. Chapter 1). Here is the law for the atomic modality. For convenience, we use the existential version.

$$\langle !\varphi \rangle \langle a \rangle \psi \leftrightarrow (\varphi \wedge \langle a \rangle \langle !\varphi \rangle \psi)$$

Interestingly, the case of iteration (and hence of future knowledge) $\langle !\varphi \rangle \langle a^* \rangle \psi$ is a bit less obvious, since we now need to make sure that we run along φ -points only.

For the recursion law, we need a system from Chapter 1, PDL with test:¹⁰²

$$\langle !\varphi \rangle \langle a^* \rangle \psi \leftrightarrow (\varphi \wedge \langle (? \varphi ; a)^* \rangle \langle !\varphi \rangle \psi)$$

But then, we really need to show that PDL as a whole has recursion laws for public announcement. This crucially involves the following technical property.

FACT 9.8 The logic PDL with test is closed under relativization.

The simple inductive proof is found at many places in the literature (cf. Harel et al. 2000). In particular, we can now state the following explicit recursion law.

FACT 9.9 $\langle !\varphi \rangle \langle \pi \rangle \psi \leftrightarrow (\varphi \wedge \langle \pi | \varphi \rangle \langle !\varphi \rangle \psi)$ is valid on process graphs.

Here $| \varphi$ is a recursive operation on PDL programs π , surrounding every occurrence of an atomic move a with tests to obtain the program $? \varphi ; a ; ? \varphi$. The effect of this transformation can be described as follows.

FACT 9.10 For any PDL program π and formula φ , and any two states s and t in $\mathbf{M} | \varphi$, we have that $s R_\pi t$ in $\mathbf{M} | \varphi$ iff $s R_{\pi | \varphi} t$ in \mathbf{M} .

This follows by a straightforward induction on PDL programs, in a simultaneous proof of the standard relativization lemma for formulas.

Knowledge We next consider the epistemic structure of forest models. The only new feature is a recursion law for the epistemic modality. This is just the standard equivalence from PAL, where we write $\langle K \rangle$ for the existential dual modality of K :

$$\langle !\varphi \rangle \langle K \rangle \psi \leftrightarrow (\varphi \wedge \langle K \rangle \langle !\varphi \rangle \psi)$$

Belief Finally, we consider the doxastic structure of forest models with plausibility orderings. The relevant law in this case is one from Chapter 7 for belief change under hard information, stated here for the modality of safe belief:

$$\langle !\varphi \rangle \langle \leq \rangle \psi \leftrightarrow (\varphi \wedge \langle \leq \rangle \langle !\varphi \rangle \psi)$$

This concludes our discussion of all relevant recursion laws for public announcement of facts about a game. A completeness proof clinching Theorem 9.4 now follows on the pattern described in Chapter 7. ■

¹⁰² This is similar to the move to the system E-PDL in van Benthem et al. (2006a).

Thus, updating forest models is an application of standard techniques.¹⁰³

Strategies As we noted in Chapter 4, PDL has the further virtue of explicitly defining strategies for players as programs. Hence the above analysis of PDL programs under public announcement also yields recursion laws for the game modalities $\{\sigma\}\psi$ of Chapters 1, 4, and 11 defined as saying that following strategy σ forces only outcomes in the game that satisfy ψ .¹⁰⁴ This leads to the following result for an extended logic PDL + PAL adding public announcements to PDL.

THEOREM 9.5 PDL + PAL is axiomatized by combining their separate laws while adding the following recursion law: $[\!|\varphi|\!] \{\sigma\} \psi \leftrightarrow (\varphi \rightarrow \{\sigma|\varphi\} [\!|\varphi|\!] \psi)$.

Proof One can use Fact 9.10 on relativizing PDL formulas and programs.¹⁰⁵ ■

Strategies involving knowledge In the current information-oriented setting, an interesting kind of strategies are the knowledge programs of Chapter 3, where test conditions have to be known to be true or false by the agent. These programs defined uniform strategies in imperfect information games. How do such programs interfere with getting more information? The above logic PDL + PAL will tell us, but the result is not always what one might expect.

EXAMPLE 9.3 Pitfalls of knowledge-based strategies

One might think that learning more, say by a reliable public announcement, should not affect the effects of a knowledge program. But this is not correct. Consider the knowledge program IF Kp THEN a ELSE b in a model where you do not know if p is the case. It tells you to do b . Now suppose you learn that p is the case, through an announcement $!p$. The knowledge program now switches its recommendation to doing a , which may in fact be disastrous compared to b . ■

103 A similar analysis applies to other temporal languages. A recursion law for the earlier branch modality $\exists G$ of Chapter 6 is as follows: $\langle !\varphi \rangle \exists G \psi \leftrightarrow (\varphi \wedge \exists G \langle !\varphi \rangle \psi)$.

104 As with Backward Induction in Chapter 8, we need converse action modalities to define rationality, but our analysis easily extends to PDL with a converse operator.

105 It may be a bit disappointing to see what the preceding result does. The recursion law derives what old plan we should have had in an original game model \mathbf{G} to run a given plan σ in the new model $\mathbf{G}|\varphi$. But more interesting issues are just the other way around. Let a player have a plan σ in \mathbf{G} that guarantees some intended effect φ . Now \mathbf{G} changes to a new \mathbf{G}' . How should σ then be revised to get some related effect φ' in \mathbf{G}' ? This seems much harder, as we noted in our discussion of understanding a strategy in Chapter 4. We will encounter similar issues below in our discussion of game change.

Other interesting questions arise when we consider strategies such as the one for Backward Induction that also involves preferences. We will return to this particular issue in Section 9.9.

Conclusion We have shown how forest models support the information dynamics of Chapter 7, allowing us a much richer account of events that can happen during play, from playing moves under uncertainty to receiving extra information beyond observed moves. However, the latter dynamics involved a radical step. Instead of viewing forest models as complete records of everything that has taken place, we now also use them as local states that can be modified when events happen that go beyond the official definition of the game. We will return to this contrast at the end of this chapter.

9.6 Forest murmurs: Soft information during play

The forest dynamics of the preceding section is easily extended to other types of event, since much more can happen in games than just getting hard information.

For a start, continuing with the techniques of Chapter 7, it is easy to add updates with soft information and plausibility change. As we have seen in Chapter 8 on Backward Induction, radical upgrades $\uparrow\varphi$ may play an important role, and these can come from many sources.

Soft triggers In fact, we usually take information coming to us as soft, unless we trust the source absolutely. This holds for imperfect information games where we can have beliefs about moves that were played on the basis of extra triggers that are suggestive but not conclusive, such as seeing a player draw a card from the stack with a happy smile. But even in public settings like an ordinary conversation, one must take careful note of what is said; but it would be foolish to burn all the bridges of alternative truths.

Complete logics We will not spell out the logic of radical upgrade over forest models, since all the ingredients are in place. Suffice it to say that we need recursion principles for the same structures as above: pure action, knowledge, and belief. For pure action, simple commutations suffice, since plausibility change does not eliminate worlds, and does not affect available actions or epistemic links. There will be axioms showing how pure action affects belief, but these are exactly the same as in Chapter 7, since the models there were fully general.

Forward Induction The issue remains of what soft updates make best sense in play of a game. A case in point is Forward Induction, a way of understanding games that was raised in Chapter 6, and that will be discussed more thoroughly in Chapter 10. Unlike Backward Induction, Forward Induction combines information from two sources: (a) observing the past of the game as played, and (b) analyzing its remaining future part. Note that the way of observing the past need not be a neutral recording of facts, as in our events $!e$ of Section 9.4. There may be more sophisticated intensionally loaded ways, such as

- (a) e was played intentionally, on the basis of rationality in beliefs.
- (b) e was played by mistake, by deviating from Backward Induction.

When we observe a move e , then, taking e to be rational gives information about the active player i 's beliefs: they are such as to make e rational in beliefs. Now we can still view this as public announcement, be it of a more informative statement:

! “move e is rational for i in beliefs.”

But since the coloring of the observation, say by rationality, may only be a hypothesis on our part, we may not want to use hard announcements, but the soft information of the radical upgrade \uparrow . This issue will return in Chapter 10.¹⁰⁶

9.7 Preference change

Many further dynamic events makes sense in games. As we have noted in Chapter 7, rational agency does not just consist of processing information and adjusting knowledge and belief. It also involves maintaining a certain harmony between informational states and agents' preferences, goals, and intentions. Therefore, it makes sense to also study preference change. Triggers for such changes can be diverse. We may obey a command or take a suggestion from some authority, establishing a

¹⁰⁶ Using \uparrow has the additional virtue that we can now make sense of any move, even those that are not rationalizable. A radical upgrade for rationality in beliefs will put those worlds on top where the latter property holds, but when there are no such worlds, it will leave the plausibility order the same. What happens in that case is just a bare observation of the move. Of course, this minimal procedure does not address the issue of how to solve conflicts in our interpretation of behavior, which may involve further updates in terms of changing preferences (see below).

preference where we were indifferent before; we may undergo a spontaneous preference change such as falling in or out of love; or we may adjust our preferences post facto, as in La Fontaine’s well-known story of the fox and the sour grapes. We have already seen how to design dynamic logics with recursion axioms for events of preference change that modify betterness ordering between worlds, working on a close analogy with the earlier plausibility changes in models for beliefs. Such logics have been studied in Girard (2008), and especially Liu (2011), to which we refer for details and applications.

Games are a typical instance of the balance between information and preference. Therefore, dynamifying their information content has a natural counterpart in dynamifying their preference structure. There are two kinds of preference change that make sense for games, or models for games. One is the realistic phenomenon of *changing goals*. Intuitively, we often do not play a game for any numerical payoff. Rather, we try to achieve certain qualitative goals, such as winning, or much more refined aims.¹⁰⁷ Things may happen that change players’ goals as a game proceeds, and when this structure changes, preferences over outcomes have to be adjusted.¹⁰⁸

Deontic views Related to this scenario is the natural connection between preferences in games and deontic notions such as obligation and permission. Preferences may come from some moral authority, encoding what one ought to do at the current stage. Normative constraints typically change as moral authorities utter new commands, or pass new laws, and this deontic preference change, too, can be relevant to games. As we noted in Chapter 2 when discussing best action, deontic logics have been applied to games in Tamminga & Kooi (2008), Roy (2011), and in many other publications.

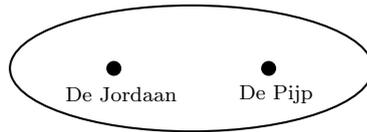
Coda: Preference change or information change? Information and evaluation are not sealed compartments. Sometimes, it is hard to separate preference change from information change. What follows is an example adapted from Liu (2011) studying the entanglement of preference and informational attitudes such as belief. For this entanglement and possible tradeoffs, see also van Benthem (2011e), and Lang & van der Torre (2008).

107 The knowledge games of Chapter 22 involve goal structure, and some intriguing results are presented there tying the syntax of goals to the solution behavior of the game.

108 Goal structure suggests the priority graphs of Andr eka et al. (2002), used in Girard (2008) and Liu (2011) to model reasons or criteria for preferences (cf. Chapter 7).

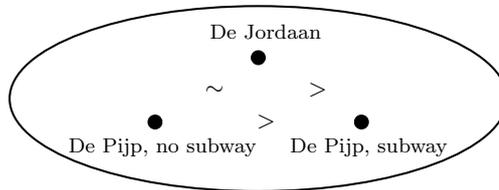
EXAMPLE 9.4 Buying a house

A potential buyer likes two houses equally, one located in the Amsterdam neighborhood De Jordaan and the other one in De Pijp. News comes out that a subway line will be built under De Pijp, endangering the house’s foundations, and the buyer comes to prefer the De Jordaan house. This starts with an initial model



with an indifference relation between the two worlds. The subway line warning triggers a preference change that keeps both worlds, but removes one \leq -link, leaving a strictly better De Jordaan house.

Alternately, however, one could describe this scenario purely informationally, in terms of a three-world model with extended options



and obvious betterness relations between them. An announcement of “subway” removes the world to the left, yielding the model we got before by upgrading. ■

This example raises issues of choosing worlds in models, the appropriate language to be used in models, and the extent to which one can pre-encode future events. No systematic comparison of the two kinds of dynamics seems to exist so far, but then, we may also view these switches in modeling just as a pleasant convenience.

9.8 Dynamics after play

Preference change also makes sense when we move from playing a game to looking back later at what has happened. Our dynamic logics apply to any activity, from prior deliberation about a game to postmortem analysis of what went wrong.

Rationalization post facto Perhaps the most effective talent of humans is not being rational in their activities, but rationalizing afterward whatever it is that

they have done. In this way, a person can even make sense of what looks like irrational behavior. If we just observe the actions of one or more players, constructing preferences on the fly, virtually any behavior can be rationalized. What follows is one of many folklore results that make more precise sense of this.

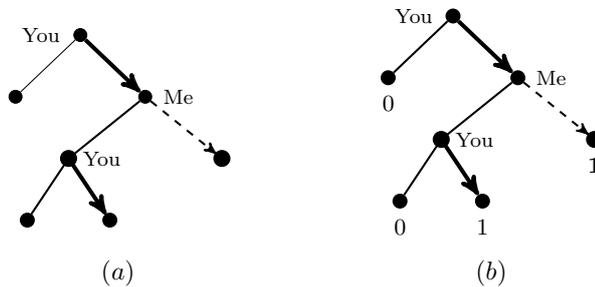
THEOREM 9.6 Any strategy against the strategy of another player with known preferences can be rationalized by assigning suitable preferences among outcomes.

Proof One algorithm works in a bottom-up fashion. Let E be the player whose moves are to be rationalized. Assume inductively that all subgames for currently available moves have already been rationalized. Now consider the actual move a made by E and its unique subgame G_a . We can make E prefer its outcome more than that of the subgame G_b for any other move b . To do so, we add a number N to all values already assigned to outcomes in G_a . With a large enough N , we can get any outcome in G_a to majorize all outcomes in other subgames G_b .

Here, crucially, adding the same number to all outcome values in G_a does not change any relative preferences in that subgame. Moving upward to turns for the other player A , nothing needs to be adjusted for E . ■

If we also assume that the player A whose preferences are given never chooses a strictly dominated move, we can even assign preferences to A to match up with Backward Induction play.

EXAMPLE 9.5 Rationalizing a game tree by stipulating preference
 Consider the sequence below, where bold arrows are your given moves, and dotted arrows are mine. Numbers at leaves indicate values postulated for you:



■

Other rationalization algorithms are explored in van Benthem (2007e), adjusting preferences or beliefs, working from the bottom up or from the top down.

Liu (2011) analyzes rationalization scenarios in terms of successive preference changes following observations of moves of a game.¹⁰⁹

Of course, many other kinds of action make sense in a post-game phase, including updates when new information changes the players’ view of what has happened.

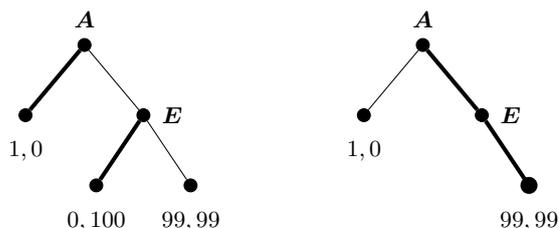
9.9 Changing the games themselves

On the road to realism about playing games, there are even more drastic scenarios. Players may not know the game they are playing: a common scenario in daily life. And if they do know the game, they may want to change it. This can happen for several reasons.

Making promises One can break the impasse of a bad Backward Induction solution by changing the game through making promises.

EXAMPLE 9.6 Promises and game change

In the following game from Chapter 8, the Nash equilibrium $(1, 0)$ can be avoided by E ’s promise not to go left. This announcement eliminates histories (we can make this binding by a fine on infractions), and a new equilibrium $(99, 99)$ results:



Intriguingly, restricting the freedom of one of the players makes both better off. ■

Game theory has sophisticated analyses of such scenarios, including so-called cheap talk (Osborne & Rubinstein 1994). In principle, these phenomena can be dealt with by our dynamic logics of information, as has been suggested in van Rooij (2004) and van Emde Boas (2011), but little has been done so far. We did see

109 Upgrades are more complex than in the above procedure: see Liu (2011) for details. Adjusting preferences also works if the beliefs of the player are given beforehand, because, as in Chapter 8, a relational strategy or a belief amount to the same thing.

an example of misleading pre-information in our discussion of safe belief in Chapter 7, a phenomenon also known from signaling games. Along a different line, Goranko & Turrini (2012) discusses pre-play negotiation in a logical framework.

Intentions and other scenarios Promises are just one instance of a more general type of event. In the dynamic logic of strategic powers in van Otterloo (2005) games change by announcing intentions or information about players’ preferences. This works at the level of the forcing languages of Part III, but one can also write such logics of game change with the above techniques. Roy (2008) uses announcing intentions to obtain simplified procedures for solving strategic games. More special, but more concrete scenarios of extraneous information flow can be found in Parikh et al. (2011) on agents manipulating the knowledge of others during play.

Other kinds of game change The borderline between changing a game or getting more information about a game is tenuous, and in principle, the above dynamic logics on forest models can handle both. Still, later chapters of this book suggest many more kinds of change that have not been studied much, and that could prove tricky. For instance, the sabotage games of Chapter 23 are about deleting moves from a given game. But as we will see there, the corresponding dynamic logic becomes undecidable, and also more complex in other ways. The same might be true with adding moves to games, adding or removing players, or changing individual powers and available coalitions. Changes in admissible sequences of moves also occur in strategy change, a topic studied by PDL techniques in Wang (2010). The general logical study of model-changing transformations relevant to games is still in its infancy (cf. van Benthem 2011e).

The importance of game change Game change may be a drastic scenario. Still, there are many reasons for taking it seriously, as we have seen already in Chapter 4. Very often, we want robust strategies that survive small changes in their current game. If plans cannot be adapted to changing circumstances, they do not seem to be of use. If a student cannot apply a method in new circumstances, the method has not really been grasped.

One more reason for robustness under change arises with the notion of rationality that is so prominent in this book. We have mentioned an analogy with theory structure in mechanics, viewing rationality as an explanatory bridge principle between various observable notions, the way that Newton’s axioms postulate forces that tie physical observables together. The main reason why this is such a powerful device in physics is the way this works, not as an ad hoc device on a case-by-case basis, but as a general way of postulating physical forces that make sense across many

different mechanical situations. By contrast, many solutions to games are fragile, falling apart as soon as the game changes.¹¹⁰

9.10 Conclusion

Chapter 8 showed how the reflection dynamics of pre-play deliberation fits with the dynamic-epistemic logics of Chapter 7. Following up on that, we took dynamic logics of informational events, and preference change, to a wide variety of mid-play processes in and around games. This involved a change from simple annotated game trees to epistemic-doxastic forest models encoding procedural information that players have about the game, and about other players. Our scenarios moved in stages. First we looked at recovering the epistemic and doxastic processes that create forest models, identifying traces satisfying the right conditions of coherence. Our results included several representation theorems. Then we moved to actual dynamics on forest models, providing a number of completeness theorems. This may be viewed as a sort of play over play, and it revealed a space of scenarios where players receive additional information, hard or soft, and even preference changes became legitimate events that can happen during play. These scenarios eventually moved into post-play evaluation, and beyond that, to changing the game itself. In all, we have shown how a wide variety of activities inside or around games can be studied in our dynamic logics.

This picture also raises new issues of its own. How do the various kinds of dynamics studied here fit together? For instance, the deliberation scenarios of Chapter 8 do not move a game forward, whereas the events in this chapter do. This is much like the reality in agency. There is the external time of new information and new events that keep changing the world, but there is also the internal time of rounds of reflection on where we stand at each moment. Somehow, we manage to make both processes happen, without getting stuck in infinite reflection, or in mindless fluttering along with the winds of change. But, just how? We will discuss this issue in Chapter 10, although we cannot promise a final resolution. There are still other general issues raised by this chapter, and we will list a few in the final section on further directions.

¹¹⁰ One might object that game change is redundant, since we can put all relevant game changes in one grand “supergame” beforehand, but as said earlier in Chapter 6, this play seems entirely unenlightening.

Despite these open problems, we have shown something concrete in this chapter, namely, how dynamic logics help turn play into a rich topic for logical analysis, far beyond the mere description of static game structure found in Part I of this book. Chapter 10 will discuss the resulting contours of a Theory of Play.

9.11 Literature

This chapter is based on van Benthem et al. (2009a), van Benthem & Dégrémont (2008), van Benthem (2007e), van Benthem (2011d), and van Benthem (2012b).

In addition, we have mentioned many relevant papers by other authors on logics that are relevant to games, and these can be found throughout the text.

9.12 Further issues

As usual, we list some issues that have not been given their due in this chapter.

Protocols This chapter and the preceding one focused on local events and their preconditions, which ignores the more general protocol structure provided by general epistemic forest models with sets of admissible infinite histories (cf. Fagin et al. 1995). We can use that structure to place global constraints on how local actions can be strung together into longer scenarios. For instance, Backward Induction in Chapter 8 used a strict protocol where only the same assertion can be made at any stage. Other protocols might restrict alternations of hard and soft information, and so on. While this can sometimes still be checked by a local counter, other protocols may need full temporal generality. We know very little about all this, witness the project for a general protocol theory in van Benthem et al. (2009a). Possible frameworks for this include dynamic logic PDL (see Wang 2010 on the related theme of strategy change from Chapter 4). Also, the theory of automata-based strategies for games that has been developed in computational logic (Apt & Grädel 2011, Grädel et al. 2002, and Ramanujam & Simon 2008) seems relevant to a better understanding of these matters (see also Chapter 18).

Histories versus stages Our dynamic logics have mainly taken a local view of temporal structure, with worlds as finite histories, i.e., stages s of some unfolding history that itself might be infinite. But we also saw a use for the temporal logics of Chapters 5 and 6 with histories h themselves as semantic objects, and evaluation of formulas taking place at pairs (h, s) . Descriptions of what games achieve often refer

to complete histories, as is clear in many chapters of this book. In Chapter 6, we also saw that the two views may bring their own takes on important notions such as belief.¹¹¹ The connection between the two perspectives remains to be understood, although we have suggested that they are tantalizingly close under suitable model transformations.

Internalizing external events, thick versus thin models In much of the dynamics of this chapter, epistemic forests change through external events. In particular, external events of public announcement simplified given models, perhaps reducing the forest to just one tree. Another approach is to internalize these external events to events that can happen inside some redesigned “supergame,” being an enlarged forest model. For instance, the protocol models of Hoshi (2009) internalize public announcements to actions inside forest models, and this has many benefits, including new protocol versions of PAL and DEL (cf. Chapter 7). On the other hand, internalizing external events blows up model size, going against the spirit of small models advocated in Chapter 6, and against our idea that complexity is best located in the dynamics, rather than in huge static models.

Clearly, there should be general transformations from one kind of model for games to another. Only with these in place would a better understanding arise of the general tradeoffs.¹¹²

Belief and expectation once more Our update formats for belief in this chapter may still be too poor when we try to really get at beliefs and expectations. An analogy may be helpful with DEL systems for updating probabilities (van Benthem et al. 2009b). The update rule mentioned in Chapter 7 turned out to need three kinds of probabilities. There were prior probabilities among worlds representing our current judgments about their relative weights; and there were observation probabilities expressing uncertainty about which event the agent actually thinks occurred. But

111 Backward Induction was a sort of intermediate scenario here. Even though we focused on its final order for endpoints, the algorithm also creates relative plausibility among stages, namely, among sibling nodes that are successors to a parent node.

112 A related issue is the extent to which update methods on forest models can work on just game trees with plausibility orders, as with Backward Induction. For instance, rough versions of Forward Induction (Chapter 10) can create plausibility order directly on trees. Start from a flat order, and consider successive nodes in the game, where moves partition all reachable outcomes. Then one can upgrade partition cells by radical upgrade for the set of outcomes that majorize, for the active player, at least one outcome for an alternative move (cf. van Benthem 2012b).

in addition to these, there were occurrence probabilities expressing what agents know about the probability that some event occurs, i.e., their knowledge of the process, encoded in probabilistic values of preconditions. The new probabilities for pairs (s, e) of an old world s and an event e then weigh all three factors, and this is important since we need to factor in how probable an event was in a given world to arrive at the right probabilistic information flow in examples such as the Monty Hall problem. Thinking in the same vein, we want a qualitative version of our update logics where we weigh plausibility of current worlds (as in our doxastic models), observation plausibility (as in our event models for priority update), but also, general plausibility reflecting the procedural information encoded in a forest model. At present, no such update systems exist.¹¹³

Connecting up different dynamics The update methods in the preceding chapters sometimes represent different takes on events. How do they interface? As a concrete example, consider Backward Induction, analyzed in Chapter 8 as a style of deliberation that created plausibility among histories of a game. But in this chapter, we worked with other belief update mechanisms. Can Backward Induction also be obtained via, say, priority upgrade of plausibility models? The answer is negative. The expectation pattern created does not satisfy the characteristic priority upgrade conditions of Preference Revelation and Preference Propagation.¹¹⁴

Even so, Backward Induction and our update mechanisms over forest models can live in harmony. We can imagine that Backward Induction has created initial expectations, and we now feed these into the real-time update process as follows. We first create an initial model whose worlds are the histories of the game, ordered

113 The challenge of finding plausibility-based update rules for procedural information, even in very simple probabilistic scenarios, is discussed in van Benthem (2012d). The difficulty is in finding qualitative analogues for the two different roles that numbers play in probability: giving degrees of strength for evidence or beliefs, but also weighing and mixing values in update.

114 In a game tree with Backward Induction-style plausibility, there can be a node x with two moves e and f where e leads to more plausible outcomes than f , while at some sister node y of x with the same moves e and f , this order reverses. For instance, assign different payoffs to the two occurrences of e and f . This stipulation also highlights the intuitive difference: Backward Induction looks ahead, while priority upgrade looks at past and present observations.

by the plausibility relation that was created.¹¹⁵ This is then the starting point for real play. In perfect information games, this consists of publicly observable moves like as discussed earlier. But with imperfect observation, further information may come in as well, either hard or soft, that can override the initial plausibility order. For instance, if we see a move that could be either an a on the Backward Induction path or an off-path move b , but with highest plausibility it is b , then we will not know where we are in the game tree any more, but the higher plausibility will be for being off-path. We will return to these interface issues in Chapter 10.¹¹⁶

Technical issues for model change Our analysis also leads to a more mathematical issue reminiscent of one raised for strategies in Chapter 4. Dynamic-epistemic logics are by no means the last word in studying model change. A standard issue from model theory fits our earlier concerns about robustness across games. When passing from one model to another where objects or facts have changed, one basic question is which statements survive such a transition. A typical example is the Los-Tarski Preservation Theorem: the first-order formulas whose truth is preserved under model extensions are precisely those definable in a purely existential syntactic quantifier form. As another example, in Chapter 1, the first-order formulas that are invariant for bisimulation were precisely those definable by a modal formula. Similar questions make sense for models and languages for games. Which assertions about games in our languages survive the changes that were relevant in this chapter? For instance, as noted in Chapter 4, it is still unknown what a Los-Tarski theorem should look like for PDL, although one has been found for the modal μ -calculus using automata techniques (D’Agostino & Hollenberg 2000).

115 The adequacy of this transformation was one of the remodeling issues raised in Chapter 6. Accordingly, we do not claim that this initial model is the optimal encoding of the Backward Induction-annotated game.

116 Similar issues arise when zooming out from Backward Induction to best action (cf. Chapters 2 and 8). Interfacing with dynamics in play involves the relativization used for common knowledge and conditional belief in Chapter 7, with recursion laws $\langle !\varphi \rangle \langle best \rangle^\alpha \psi \leftrightarrow (\varphi \wedge \langle best \rangle^{\varphi \wedge \langle !\varphi \rangle^\alpha} \langle !\varphi \rangle \psi)$. For instance, in our example of game change by a promise, let p denote all nodes except the end node with $(0, 100)$, and q just the end node with $(99, 99)$. Recomputing Backward Induction in the smaller game, we saw that q resulted. Analyzing the matching assertion $\langle !p \rangle \langle best \rangle \langle best \rangle q$ by the given law, we find that it reduces to $p \wedge \langle best \rangle^p (p \wedge \langle best \rangle^p q)$. To see that the latter is true, recursion laws no longer help, and we must understand the relativized best-move modality. This shifts the burden to understanding the conditional notion $\langle !\varphi \rangle \langle best \rangle^\alpha \psi$ plus its static pre-encoding. This refines our conjecture in Chapter 2 about axiomatizing best action.

Beyond literal preservation, translation of assertions makes sense as well, deriving modified strategies for players, or new descriptions of their powers, in games that have undergone some suitable simple definable change. No general theory seems to exist that can be applied to games without further ado.

Dynamic logic and signaling games This chapter has proposed a richer logical style of analysis for games than the usual static ones. Still, this is a program, and it remains to be applied more systematically in game theory. One obvious area for applying dynamic techniques would be the theory of signaling games (Lewis 1969, Cho & Kreps 1987, Osborne & Rubinstein 1994, Skyrms 1996, van Rooij 2004), where agents send signals about the state of the world in which a game takes place.

10 Toward a Theory of Play

It is time to pull together some threads from the preceding chapters. Chapter 7 showed how we have logical systems at our disposal for modeling actions and events that make information flow, and also for changing preferences in a concomitant stream of evaluation. In Chapters 8 and 9, we explored how this rich array of logical tools applies to various sorts of dynamic events that happen within, or alongside games. But where is all this heading? The purpose of this discursive final chapter of Part II is to combine all of these threads by thinking about a general enterprise that seems to be emerging at the interface of logic and games, which might be called a “Theory of Play.” We will not offer an established theory, but a program. Based on the topics in earlier chapters, we will discuss major issues in its design, but also critical points, and further repercussions for the logical study of social activities, gamelike or not.

10.1 Dynamics in games

Many activities happen when people engage in playing a game, with phases running from before to during and after. Chapter 8 was about pregame deliberation, and our emphasis was on procedures that create prior expectations. Taking Backward Induction as our running example, we saw how to construe deliberation as a dynamic process of iterated belief revision, working with extensive game trees expanded with plausibility relations. But this was just a case study. This initial phase contains many natural activities of deliberation and planning. Next, during actual play, many further things can happen. In Chapter 9, we saw a wide variety of events, such as playing moves, public or private observations of moves, but also events of receiving further information about the game and its players, and of

changing beliefs or preferences. Again, this list is not complete, as it all depends on the level of detail that is of interest, whether coarser or finer. For instance, many authors assume that players know their own strategies, but this amounts to assuming that a decision has already been taken, whereas, at a finer level, acts of decision themselves might be objects of study.

We have shown how all of these events can be dealt with in dynamic logics, in both hard and soft varieties, where the structures undergoing change were usually epistemic-doxastic forest models encoding procedural and social information that players have about the game and about each other. Finally, turning to the post-game phase of reflection, the same logics dealt with rationalization or other activities that take place afterward. In fact, they even dealt with more drastic events that act as game changers.

This creates a huge space of possibilities. In what follows, we try to get some focus by doing a case study, looking at how one would normally play a game, while identifying a number of issues arising that are of more general interest for the role of logic. Our case is a confrontation of Backward Induction as a style of deliberation to reasoning in actual play, with Forward Induction as an alternative style. Issues that will arise include belief revision, managing hypotheses about other agents, options in modeling update steps, the role of simple cues in normal ways of speaking about social action, and the resulting desiderata on design of logical systems. After that, we will discuss general issues in a Theory of Play, including the role of agent diversity, the tension between sophisticated description and model complexity, and possible repercussions beyond games to other fields interested in a broader theory of social action.

10.2 Problems of fit: From deliberation to actual play

Our extensive discussion in Chapter 8 may have suggested that Backward Induction is the view of games officially endorsed by logic. But this is not true. While its elegant bridge principle of rationality is appealing, and hence should not be given up lightly, exploring alternatives makes sense. Our limit scenarios of iterated hard and soft announcements $(!\varphi)^\#$ and $(\uparrow\varphi)^\#$ would work with any formula φ in our language, whether it produces the Backward Induction solution or not, and Muddy Children-style alternatives such as “iterating worries” could well be another option. In Chapter 9, our dynamic logics supported many sorts of events that can override the initial expectations created by a solution algorithm, for instance, in the form

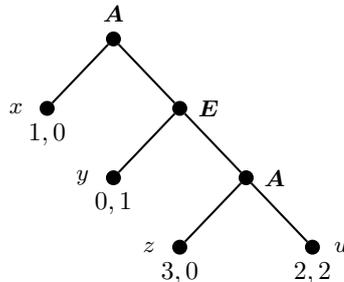
of soft updates that change the plausibility ordering. We now turn to what may be the most radical challenge.

From deliberation to actual play In moving from deliberation to reality, some well-known problems arise. Let us accept Backward Induction as a prior deliberation procedure. What about the dynamics of actual play, when these expectations are confronted with what we actually observe? One might think this is a simple matter of following the virtual moves computed in the deliberation phase. But often this makes little sense when run in this opposite direction. We expect a player to follow the path computed by Backward Induction. So, if the player does not, we must perhaps revise our beliefs, and one way of doing that is precisely having second thoughts about the player's style of deliberative reasoning.

This issue has been dramatized in the so-called paradox of Backward Induction. Why would a player who has deviated from the computed path return to Backward Induction later on (Bicchieri 1988)?

EXAMPLE 10.1 Expectations meet facts

Consider the following game, with outcomes evaluated by **A** and **E** as indicated:



Backward Induction tells us that **A** will go left at the start. So, if **A** plays *right*, what should **E** conclude from this surprising observation? **E** might still assume that **A** will play Backward Induction afterward, thinking that a mistake was made. But this is not always a plausible response. In many social scenarios, a participant **E** might rather think that the observed deviant move refutes the original style of deliberation about **A**, who may now be thought to be on to something else, maybe hoping to arrive at the outcome (2, 2). ■

In other words, although the Backward Induction strategy *bi* is easy to compute, does it make sense when one tries to interpret off-equilibrium parts of the game? The assumption that Backward Induction will prevail later on, no matter how often

we observe deviations, is the technical core of the results by Aumann (1995), and within dynamic-epistemic logic, Baltag et al. (2009), that characterize *bi*.¹¹⁷

From the game to policies of its players But clearly, this is not the only way of revising beliefs here. Player *E* could have reasonable other responses, such as

- A* is saying he wants me to go right, and I will be rewarded if I cooperate.
- A* is an automaton with a rightward tendency, and cannot act otherwise.
- A* believes that *E* will go right all the time.

We cannot tell which response is coming, unless we also know at least players’ *belief revision policies* (Stalnaker 1999, Halpern 2001). This seems right, but adding this parameter is a momentous move in the foundations of game theory. Halpern (2001) is a masterful logical analysis of how Stalnaker’s proposal undermines famous results such as Aumann’s Theorem that common knowledge of rationality implies that the Backward Induction path is played (Aumann 1995). This is true on the standard notion of rationality, but it is false on Stalnaker’s generalized view, where rationality refers to players’ beliefs about the current strategy profile, which can change as surprise moves are played.

What does all this mean for our approach? Our dynamic logics are well-suited to analyzing belief change, so technically, accommodating players’ revision policies poses no problems. Our logics are also welcoming: they do not impose any particular policy. The systems of Chapter 7 supported many update methods, of which the one generating Backward Induction is just one. Our subsequent analysis in Chapter 8 may have seemed to favor this scenario, but this was for concreteness and as a technical proof-of-concept, and we identified choice points along the road, for instance, concerning the strong uniformity assumptions underlying the *bi* strategy.

Broader issues in social action But these issues are not just technical. They reflect real challenges for a logic of social action. Our expectations may be based on prior deliberation, including scenarios that we think will not occur. But what if the unexpected happens? A move considered hypothetically may impact us quite differently once it has occurred: cold feet are ubiquitous in social life. How can a priori styles of deliberation and actual play of a game that is constantly being updated with observed events be in harmony? There are often deviations from this harmony in practice, but it is definitely an interesting ideal.

¹¹⁷ Baltag et al. (2009) call this attitude an incurable belief in rationality.

Player types Belief revision policies need not be our only focus in studying these phenomena. Social action involves many kinds of hypothesis about other agents. For example, an agent may perform an action based on assumptions about the memory capacity of other agents or their learning behavior. This point is acknowledged in game theory. As we saw in Chapter 6, “type spaces” are meant to encode all hypotheses that players might have about others. But our problem is the distance from reality, where we get by with simpler views of the social scenario we are in. Hence, we have worked with simpler models for games, describing the dynamics of relevant events instead of having things be prepackaged.

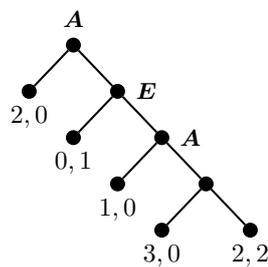
To make these new perspectives more concrete, we will now discuss a case study of social reasoning, pursuing a simple view that lends itself to dynamic logical analysis. Our treatment will be light, and we do not offer a final proposal. After the case study, we take stock of some general features of a logical Theory of Play.

10.3 Forming beliefs from observations

Let us discuss the earlier example in a bit more detail, focusing on its update aspects. In the bottom-up Backward Induction computations of Chapter 8, we omitted the history of play: it did not matter for our expectations how we arrived at the present node. But just think about how you yourself operate in social scenarios: the past is normally informative, and we do need to factor it in.

EXAMPLE 10.2 Interpreting a game

Here is a simple example, varying on an earlier game:

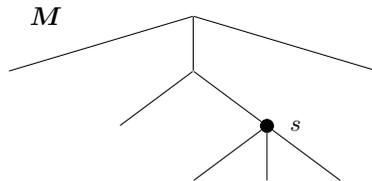


Suppose that **A** chooses *right* at the start. Assuming that **A** is rational-in-beliefs, this informs **E** about **A**'s beliefs and/or intentions. Clearly, **A** does not expect **E** to go left at the first turn, because then, playing *left* at the start would have been better. For the same reason, **A** does not intend to go *left* after that at the second

turn. And we may furthermore assume that A believes that E will go *right* at the end, as that is the only way the opening move makes sense. All this might induce E to go *right* at the first turn, although one hesitates to predict what move will be chosen at the end. ■

The point here is not that we have one simple rule replacing Backward Induction. It is rather that the past is informative, telling us which choices players made or avoided in coming here. Our observations and our expectations work together. Now we do not have one unique way of doing this, but there are clearly intuitive scenarios reflecting our own practice.

Games with a history The change needed is easily pictured. We now look at games M with a distinguished point s indicating how far actual play has progressed:



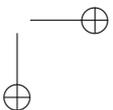
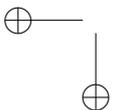
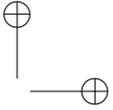
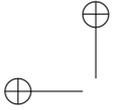
Thus, at least in games of perfect information, at the actual stage, we know what has happened, and players can let their behavior depend on a mixture of two things:

- (a) What players have done so far in the context of the larger game.
- (b) What the structure of the remaining game looks like.

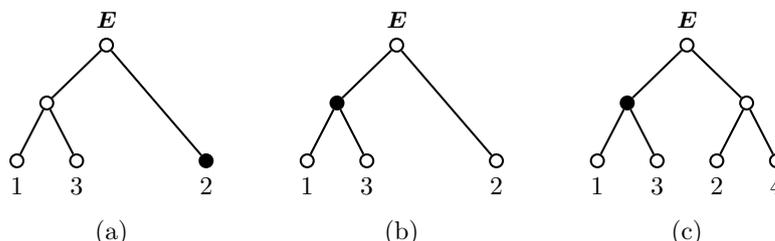
We will see more complex models later, but for now, we will discuss simple scenarios.

Ways of taking observed behavior How will players change their expectations as the black dot moves along a game? Readers can hone their intuition on simple cases of making decisions.

EXAMPLE 10.3 Making sense of decisions
We picture two simple scenarios:



To the left we see the end of a basic rational decision. To the right we see a “stupid move,” probably regretted by E once made.¹¹⁸ Here are a few more complex cases:



The play observed in game (a) may be considered rational by ascribing a belief to E that choosing left would have resulted in outcome 1. Game (b) may be rationalized by ascribing a belief to E that the game will now reach 3. Finally, game (c) suggests that E thinks that 3 will be reached, while 2 would have been reached if the initial move had been to the right. ■

Rationalizing There are many options for making sense of observed behavior in the preceding example. Actual moves may be considered to be mistakes, or even as signs of stupidity. They may also be taken to be smart, but how smart depends on assumptions. Let us discuss one natural tendency when making sense of what others do. We stay close to rationality, and only drop that assumption about others when forced:

Rationalizing By playing a publicly observable move, a player gives information about beliefs. These beliefs rule out that the player’s actual move is strictly dominated in beliefs.

This will only work if the player is minimally rational, not choosing a move that is strictly dominated under all circumstances, i.e., under every possible continuation of the game.¹¹⁹ Only if rationalization does not work might we consider stupidity or some other obstacle. This suggests a ladder of interpretative hypotheses, where we move to the next step only when forced. But for now, let us stay with rationalization.

118 Regret seems central to social life, although it may work best with iterated games.

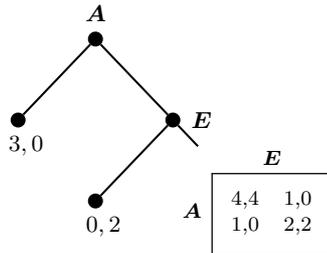
119 This “weak rationality” avoids stupid things. One could also make the stronger assumption of “strong rationality,” where the agent thinks the chosen move is best. Chapter 12 considers counterparts of both notions for strategic form games.

Implementing rationalization still depends on additional assumptions about the belief structure of the agent. In the presence of minimal rationality, and beliefs allowing for ties, one way of rationalizing is simply to assume that the agent considers all continuations to be equally plausible. No observed move can then be strictly dominated in beliefs. But assuming that agents have one unique most plausible history in mind, more information comes out of an observation. Unique belief plus strong rationality were in fact reasons for suggesting that in the above game (c), agent E believes that 3 will be reached, and that 2 would have been reached after going right. Fixing stipulations, we get various algorithms, all proceeding on the principle that moves reveal beliefs about the future.¹²⁰

Forward Induction Scenarios such as the above are sometimes called Forward Induction (cf. Brandenburger 2007), suggesting a simple computational change to the Backward Induction algorithm, or relatives for strategic games such as Iterated Removal of Weakly Dominated Strategies, that will remedy the earlier-noted deficiencies. Whether this remedy is possible or not, we now explore a relevant question about switching algorithms.

Can dropping Backward Induction be for the best? Sometimes, dropping Backward Induction may be advisable for rational players.

EXAMPLE 10.4 A Forward Induction scenario
The following game is adapted from Perea (2011):



In the matrix game, no move dominates any other. Hence, E considers all outcomes possible in Backward Induction. Then going left is safer than going right,

¹²⁰ Incidentally, the beliefs of a player E do double duty in this setting. At a turn for the other player A , they are real beliefs about what A will do. But with a turn for E , they are more like intentions.

and hence A should go left at the start. But if E rationalizes, and sees A go right, there is extra information at the next point: A expects to do better than 3, which can only happen by playing *up* in the matrix game. This tells E to go there, too, and play *left*, doing better than the Backward Induction move giving E just 2. ■

Top-level views and fuzzy endings? While this scenario is interesting, it needs a convincing interpretation for the final matrix game. Often we do not know the complete game, or it is beyond our powers to represent it. We only know some top-level structure. The matrix game of imperfect information then gives a rough image of what might happen afterward. This goes against the spirit of analyzing games all the way to rock bottom, the way we did in Chapters 2 and 8. But it is like solution algorithms in AI that work with some top-level game analysis plus general heuristic values for unanalyzed later parts of the game (cf. Schaeffer & van den Herik 2002). This is also closer to the way we often reason in real life.¹²¹

Coda: Extending rationality Our discussion is not meant to suggest that rationalization in the sense of Forward Induction is the only alternative to Backward Induction. An interesting alternative is minimizing regret, proposed in Halpern & Pass (2012), and studied using dynamic-epistemic logic in Cui (2012). Yet another view is suggested in van Benthem (2007d), based on the social phenomenon of returning favors, compensating agents for risks they have run for improving the payoffs for both of us. This points at perhaps the most general view to emerge from our discussion in this section, that of weighing both future and past. A player acts in a “responsibly rational” way by taking care of that player’s own future interests while giving past interactions with the opponent their proper due. Cooperation deserves consideration; lack of cooperation justifies neglect. This is how most of us navigate through life, and it would make sense in many games.

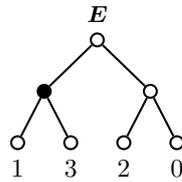
10.4 Logical aspects: Models and update actions

We started by analyzing Backward Induction as a process of deliberation prior to playing a game. We have now discussed how to analyze a game as it is being played, making a case for considering alternatives using information from the past. We now explore a few further logical aspects of this setting, tying in with earlier chapters.

¹²¹ The general algorithm in Perea (2011) raises further logical issues that we forego here.

Choosing the models Models in Section 10.3 were pointed trees that mark where play currently stands. The pointed forest models of Chapters 5, 6, and 9 were like this, although more general in allowing for variation in what players know or believe, as encoded in epistemic and plausibility relations between nodes. This is the generality that we need. For instance, to describe the earlier rationalization procedure, we need general forest models, since beliefs need not have a simple *bi*-style encoding as best moves. Quite different beliefs for an agent may rationalize an observed move, and these need not have a weakest common case.

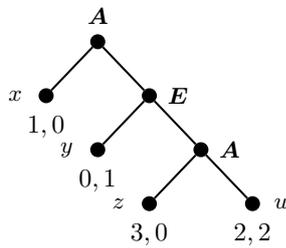
EXAMPLE 10.5 Incompatible hypotheses about belief
 Consider the following tree. There are two incompatible ways of rationalizing player *E*'s left move arriving at the black dot:



One option assumes that *E* expects 3 on the left, and any outcome on the right. Another option lets *E* expect any outcome on the left, and 0 on the right. ■

Updating thinner or thicker models We also have different update scenarios for the earlier games, depending on whether we choose thinner or thicker models.

EXAMPLE 10.6 Expectation meets facts, revisited
 Consider this earlier game from Example 10.1 in Section 10.2:

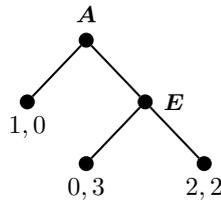


One way of interpreting the *right* moves in the game is as follows. We start with equiplausibility for all branches for all players. *A*'s first *right* move triggers an

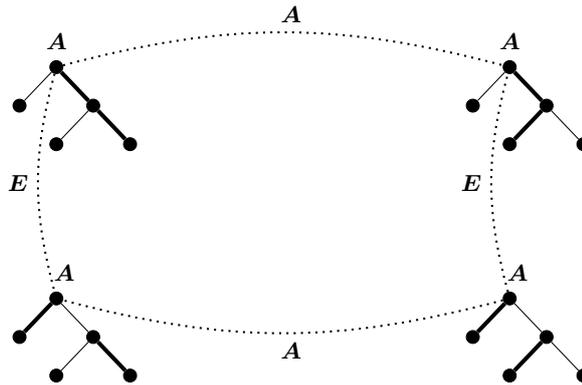
upgrade, making histories RRR , RRL more plausible than RL in the ordering \leq_A . Next, E 's right move makes RRR more plausible in \leq_E than RRL and RL . ■

However, sometimes a simple plausibility shift may not do the job, and we will have to update more complex models. This may be seen in the following example, a variation on a recurrent illustration throughout this book.

EXAMPLE 10.7 Updating thicker models
Consider the following game:



The four possible strategy profiles pre-encode all possible response patterns. These form a forest with epistemic links as indicated, marking a moment when players have decided on their own strategy, but do not yet know what the other will do. We only indicate the top-level links. Lower links will match corresponding lower nodes, reflecting our assumptions of perfect recall and public observation.¹²²



¹²² Epistemic forests also allow finer distinctions, such as assuming that after the moment of decision, players know their next move, but not their entire strategy yet.

Let the top left tree be the actual one, which means that \mathbf{A} plays *right*. Even then, we still have a choice of update. One view is that moves are played according to players’ intentions. That is, an event e takes place with the precondition, “ e is a move prescribed by the active player’s strategy,” and the update would leave only the two topmost models, making \mathbf{E} , but not \mathbf{A} , know how the game will end. But we can also assume that players make mistakes, with a weaker precondition “ e is a move that is available to the active player.” Then e may also have been played by mistake, and all four subtrees remain, shifting the current stage of the model, but nothing else. ■

This model can be made more complex to also allow for players’ expectations, and their belief changes under public information.

Options galore What we see here is a tradeoff between plausibility upgrade on trees and pre-encoding global information in game models about agent behavior, but then using simpler updates such as public announcement.¹²³ In addition to this choice of models, there is a variety of choices for updates interpreting events. For instance, we can decide to take moves at face value or with stronger intentional loading, and we can model decision steps explicitly, or leave them implicit. Also, forces can differ, in that rationalization steps need not be public announcements, but could be soft upgrades with rationality, as in the second scenario for Backward Induction in Chapter 8. Likewise, as for the models involved, we can update close to trees, or simplify updates in richer models pre-encoding global information about agent behavior. This variety seems true to the many options that humans have for making sense of behavior.

Now we move to a technical issue that goes slightly beyond our earlier logics.

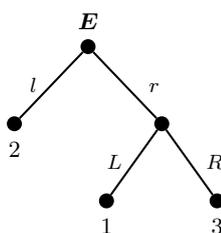
Ternary plausibility There is also a generalization in the air here of the earlier models of Chapters 7, 8, and 9. While Backward Induction created one uniform binary plausibility relation $x \leq y$ among histories x, y , our discussion of Forward Induction suggested a ternary plausibility relation $\leq_s xy$ (cf. van Benthem 2004c) where the ordering may depend on the current vantage point s . This ternary ordering allows for differences between what players expect hypothetically if another move had been played than the actual one (“that would have been stupid”) versus how they would feel if that other move were actually played.

123 A similar tradeoff occurred in Chapter 9 for preference change and information change.

Backward Induction had no distinctions based on current viewpoints in its plausibility ordering, but more general procedures of rationalization need not produce expectations that match up across a game. Our rationalization procedure used off-path expectations in the game as a contrast, allowing us to get more information about relevant beliefs in the future of our current path.¹²⁴ Recall Example 10.2, which we now explore more fully.

EXAMPLE 10.8 Crossing expectations, revisited

This time, read numbers, not as utility values, but as degrees of plausibility:



This violates the uniform node-compatible plausibility order for Backward Induction (cf. Chapter 8) that makes one of the moves l , r more plausible than the other, while all their outcomes follow this decision. Yet it is easy to find scenarios where the order depicted is natural, and general models seem more appropriate.¹²⁵ ■

DIGRESSION Here is a common objection to this move. Is the technical move to ternary orders coherent? Surely, a current plausibility order already determines what is plausible at other nodes by counterfactual reasoning, or does it? A player chooses move a , and says “if I had chosen move b , then the following would have happened.” Where is the need for making the latter reasoning dependent on actually being at the node reached by playing b ? This objection seems confused. If a player has in fact chosen move a , then that player considers choosing b a mistake, and the counterfactual has a hidden assumption of being in a bad place. But if the

¹²⁴ There need not even be consistency going down a path: expectations may flip when a player makes a nonrationalizable move. Technically speaking, Backward Induction might be the only uniform consistent algorithm creating expectations.

¹²⁵ The uniform plausibility relations produced by the Backward Induction algorithm are then a subclass defined by a special axiom in the doxastic-temporal language, whose details we do not spell out here.

player has actually chosen b , judgments will proceed on the assumption that this move was chosen intentionally. Thus, a counterfactual cannot be taken at face value. It can only be evaluated if we make our assumptions explicit about how the move was played: intentionally, by mistake, and so on. These assumptions generate different plausibility orderings, and so the third argument reappears. With Backward Induction, the built-in assumption is that all off-path moves are mistakes.

Dynamic logic over ternary models Ternary ordering relations are a standard tool in conditional logic (Lewis 1973), and ternary versions of our dynamic logics of Chapter 7 were applied in Holliday (2012) to epistemological views that explain knowledge in terms of counterfactual beliefs. Such logics would have to be adjusted to our earlier topics, for instance, tying the limit scenarios of Chapters 7 and 8 to standard fixed point logics such as LFP(FO) or IFP(FO) that easily tolerate ternary predicates.¹²⁶

Update with agent types Our earlier scenarios largely used local updates, where views of the game change as certain events are observed and given a particular interpretation by players. But as we have suggested, there is also room for an additional global view of the type of agent we are dealing with. Backward Induction presupposes a particular kind of future-driven rational agent, while rationalization modified this assumption by allowing agents other beliefs, revealed through their behavior. Still further agent types emerge when such rationalization of behavior fails, perhaps moving to hypotheses of impaired rationality.¹²⁷ While these may seem to be pessimistic ladders of agent types descending into inanity, there are also optimistic views, starting with an assumption that the player is a simple machine, and moving up to more complex process views of others as needed. Modeling such ladders fits well with our earlier models, but it requires a more complex structure of update moves than what we have investigated.

EXAMPLE 10.9 Games with different agent types

Assume that players do not know whether the game they are entering is cooperative

126 Our discussion also suggests more radical challenges to the logics used in this book. For instance, one case of Forward Induction looked at extensive games that end in strategic matrix games, or more simply, in open-ended games whose structure we do not have available. What happens to our logical analysis when we allow this sort of hybrid model?

127 Similar ladders occur in social life, where we start by assuming that people are friendly and reasonable, and give up these illusions only when forced to by new facts at hand.

or competitive, a typical situation in daily life. But they do know that there are only two types of player: competitive (playing Backward Induction) and cooperative (striving for a best cooperative outcome). Of course, what the latter means remains to be defined in more detail, but even now, we can see how this may drastically change earlier update scenarios. If we only have these two types, then one observation may tell us the type of the opposing player and all of the future behavior, and the game may reach a stable situation after a few moves. ■

While this scenario looks attractive, and while working with a small set of possible types sounds realistic, it also has an almost static flavor, in that opposing players of fixed types are predictable automata. But there is an interesting asymmetry in much social reasoning. Often, players may see themselves as unique and infinitely flexible, while opponents fall into general types of behavior encountered before, as in Oscar Wilde’s famous snub: “I don’t know you, but I know the type.” The same asymmetry can be seen in many logic puzzles (cf. Liu & Wang 2013): the inhabitants of Smullyan islands are predestined Liars or Truth Tellers, but the visitor is a flexible human trying to design questions so as to detect the inhabitants’ type and then profit from it.¹²⁸ The most interesting types seem those allowing for sufficient variation in behavior.

I will not develop any logic for these scenarios, but see Wang (2010) and Ramanujam (2011) for some computational logic-inspired analyses of agent types. What might also be of interest is extending standard characterization results in the foundations of game theory. Aumann’s Theorem says that if it is common knowledge that all agents are of one particular type, that of standard rationality, then the Backward Induction path is played. What if we relax the assumption, and assume as common knowledge that each agent is either competitive in this sense or cooperative? Could there be an extended logical theory of games in terms of natural families of agent types?

10.5 Theory of Play

We now draw a conclusion from the preceding considerations. There is no unique way of defining the best action in a game. The missing ingredient is information about the types of agent we are interacting with. The structure of a game by itself

¹²⁸ Logic puzzles would get a lot harder if all participants had flexible types.

does not provide this information, unless we make strong uniformity assumptions. We need more input.

The term coined in van Benthem et al. (2011) is “Theory of Play.” To make sense of what happens in a game, we must combine information about game structure plus the agents in play. Game theory allows each player to have personal preferences, but on top of that, say, many solution procedures assume uniformity on how players think and act. But we need much more variety: in cognitive or computational limitations, belief revision policies, and other relevant features. Further, what is true for games is true for social scenarios in general, the players matter. Actually, there are two aspects to this extension: we need to know about the actual play, and also about the players involved. The two are intertwined, but they represent different dimensions. Logic can help with bringing out the variety of scenarios and reasoning styles that go with this.

Of course, this idea is not totally new in the literature. Many developments in game theory tend in this direction, and the same is true for much work on game logics. Still, the current phrasing seems useful as a way of highlighting what is involved. It may also be important here to reiterate a point from our Introduction. Theory of Play is not something that is going to cure the current ills of game theory, or of logic. Rather, it is a joint offspring of ideas from these two fields (and others, such as computer science). Children are often a stronger bond in relationships than are personal transformations.

At present, there is no well-developed Theory of Play but only interesting bits and pieces that might help create one, among which are the rich set of logical tools developed in Part II of this book. This chapter concludes in a discussion of a few major challenges and potential benefits of the enterprise.

10.6 Locating the necessary diversity

One central concern is where to locate the variety that is needed in a Theory of Play. Possibilities are vast: the multi-agent system of the players, or just their strategies, or perhaps their repertoire of interpreting behavior by others, or at least reasoning about it, letting variety be in the eye of the beholder.

Individual events versus general types One can even ask whether we need agents at all, since they are a temporal entity with behavior persisting over time that goes far beyond the few events observed so far. Why not stick to what we observe

locally and update based on that? The tension between individual events and postulating general types beyond these seems typical of human language and reasoning. We tend to see each other in generic terms, routinely using adjectives such as friendly or hostile that package a whole style of behavior over time. Cognitively, a genuinely isolated one night stand is about as rare as a free lunch. Presumably this tendency toward instantaneous generalization has a cognitive value, since just sticking to the facts would turn us into mindless signal-recording devices.¹²⁹ We therefore turn to agents.

Taxonomy of players There are huge spaces of possible hypotheses about players, but real understanding involves finding smaller manageable sets of relevant options. For instance, Liu (2009) has a nice map of agent diversity from the standpoint of dynamic-epistemic logic. One dimension is the processing properties of agents: what are their powers of memory, observation, and even of inference?¹³⁰ Another dimension is the update policies of agents: how will they revise their beliefs, or more generally, what learning methods do they follow? A third dimension might be called “balance types” between information and evaluation: agents can be more optimistic or pessimistic in pursuing their goals, and so on. Finally, also relevant might be social types such as whether players are more competitive or more cooperative, as discussed earlier.

Sophisticated versus simple strategies Agent types are one way of doing things. We might also just consider strategies as partners in interaction. Taxonomy of agents then gives way to taxonomy of strategies. Here lies a challenge to the logical approach in this book with its emphasis on ever greater sophistication in epistemic reasoning. Many studies of social behavior show that simple rules often work best (cf. Axelrod 1984, Gigerenzer et al. 1999). A player may be a sophisticated intellectual full of theory of mind, but perhaps the only thing that matters right now is whether the player is following a simple strategy of Tit for Tat. We are far from a general understanding of when simple strategies suffice, and when sophisticated reasoning about others is really needed.

129 This dismissal of individual events may be too hasty, and there is an intermediate option. Ramanujam (2011) is an intriguing exploration of a space of agent types that grow over time in a social process.

130 Along this line of thought, off-path behavior against Backward Induction may indicate the presence of another reasoning style by the relevant player.

Diversity in logics Diversity also abides in our logics. The logical dynamics of Chapter 7 highlighted the diversity of observational access or plausibility shifts by various update mechanisms, greatly extending standard views of logical agents. And even the presuppositions of our logics can be varied. The examples explored in Chapter 9 involved games with dynamic-epistemic update. In line with agent diversity, there were also update rules for memory-bounded agents whose epistemic-temporal forest patterns could be determined.

In other words, our approach is diversity-tolerant. But could it be too much so?

10.7 Some objections

There are certainly objections to what we are proposing. We address them directly.

Messiness Theory of Play comes at the cost of a large space of hypotheses about agents, making models quite complex. How can this explosion be kept in hand?

This objection has a good deal of merit, and it can be a salutary force for keeping things simple. For instance, our study of rationalization used complex updates over complex models from Chapter 9. Perhaps we should instead look for simpler alternatives. In particular, interpreting moves in social scenarios may involve just a small number of ways that are common in practice. In addition to public observations *!e*, these might be “uptake” acts considered earlier such as:

- e* was played intentionally, on the basis of rationality in beliefs.
- e* was played by mistake, by deviating from Backward Induction.¹³¹

Losing the appeal of uniformity Uniformity assumptions such as those embodied in Backward Induction are not just a simplistic modeling view to be replaced by sophisticated diversity. They also represent attractive intuitions of treating people equally, while reflecting a crucial intelligent ability of being able to put oneself in someone else’s place (cf. van Benthem 2011d). Moreover, much cognitive behavior is held in place by forms of resonance between similar minds, so we should not give up uniformity lightly.

¹³¹ Moreover, given the fact that these additional features may be hypotheses on our part, we may want to use these either as hard information or soft information.

Certainly this objection has some force, but perhaps resonance occurs at some higher general levels (for instance, by agreeing to play a game at all), while diversity reigns at more specific levels.

Understanding too much The apparatus developed here can model virtually any behavior. Where is the normative force that is crucial to criticizing behavior, another aspect of taking people seriously, rather than letting them stew in their own juices?

There is no clear response to this quite reasonable objection. Social life is a delicate matter of balance, and perhaps, so is its logic.

10.8 Living with diversity, linguistic cues, and logical design

We end with some more constructive thoughts on the issues raised so far for a viable Theory of Play.

Using information that we have The preceding objections may make things look more complex than they really are. There are also forces that strike a blow for simplicity. Normally, we do not have to produce hypotheses out of the blue. Our expectations about people are based on earlier experience, so we do not enter social scenarios as a tabula rasa. And even though puzzles in the literature seem lifted out of context, often the concrete description of the scenario can be mined for agent types.¹³²

Social life and language Coping with diversity is a fact of successful social life that takes several sources of tension in stride, such as the earlier division between thinking in terms of types or just responding to individual events. While this may not be totally convincing (is social life really so successful?), and while appeals to the facts are often the last resort of desperate theorists, looking at empirical evidence may be important at this stage of theorizing. For instance, one underused resource is our natural language. We have a rich linguistic repertoire for talking about individual behavior and social interaction, of which only a tiny part has been studied by logicians. Just think of all of the terms like regret, doubt, hope, reward, revenge, and so on, that structure our lives, while having a clear connection

132 For instance, a famous probabilistic puzzle like the Monty Hall problem can only be solved if we know which protocol the host is following (Halpern 2003b, van Benthem et al. 2009b). Such information can usually be found in the statement of the scenario.

with the balance of information, evaluation, and action that is so crucial to games. This repertoire seems to serve us well, so it might provide a sort of anchor to the logic of social interaction.¹³³

Finding unity in all the right places Moving to more technical perspectives, one can also wonder how much unity of methods is needed, and where it resides. For instance, while dynamic logics allowing agent variety may get complex, a counteracting force is “redesign.” Consider the powers of observation. One might write different logics for all sorts of agents with varying access to what is happening. But dynamic-epistemic logic packs all of this variety into one relevant event model, and then describes one mother logic for updating with these.¹³⁴ The same was true for belief revision. Prima facie, it dissolves into many update policies, with a resulting jungle of logical systems (see van Benthem 2007c on complete dynamic logics for many policies). But Baltag & Smets (2008) let event models encode the variety again, leaving only one rule of priority update obeying a simple complete set of axioms that we saw in Chapters 7 and 9, be it at an intuitive price of having more abstract signal events. And the discussion on best logical formats for social scenarios continues (cf. Girard et al. 2012).

Thus, Theory of Play should acknowledge diversity, while taking full advantage of all available cues, and letting logic do its usual job of abstraction and idealization.

10.9 Connections and possible repercussions

Agent diversity and theory of play make sense beyond games. For instance, computer science has a large body of results on what can be achieved by different kinds of strategies (Chapter 18 surveys some results, generalizing the work on logic games in Part IV). Likewise, behavior in cognitive experiments illustrates the earlier mismatch between deliberative rationality and actual play, because preferences may change in the heat of battle.¹³⁵ Our theory should be informed by all of this.

¹³³ We have based the logics of this book on calm beliefs and preferences, but what if we base them on the warmer hopes and fears that inform our real decisions?

¹³⁴ Admittedly, finding that mother logic can be highly non-trivial, witness the discussion in Chapter 7 on recursion laws for DEL with common knowledge.

¹³⁵ Compare this to McClure (2011) on behavior in auctions deviating from game theory.

Theory of Play may also affect fields such as philosophy that are packed with uniformity assumptions, often based on philosophical intuitions that serve as a uniformizing standard. What happens when we question these assumptions? What is fair play in ethics given the undeniable diversity of agents? Are the usual Kantian ideas about all of us reasoning in the same way the greatest justice, or the greatest form of injustice? Or consider epistemology: what is rationality? Is it doing the best by your own lights, no matter how dim? Similar points apply to the philosophy of language, where the usual models of meaning involve uniform language users that belie the variety of actual communication. There is a tension between the lofty impartiality of uniformity assumptions and the humanity of diversity views, but in any case, both deserve a hearing in our logics.

Theory of Play might even reach logic itself. What about a Theory of Inference describing human or computational agents engaging in deduction and other activities, and their different styles of doing so? Say, different kinds of automata engaging in proof search, or competing in logic games? Can logic get closer to actual reasoning if we relax its standard uniformity assumptions about agents that remain implicit? Might this lead to a new understanding of existing formal systems, when we study them in use?

10.10 Conclusion

This chapter has drawn together the threads of Chapter 7, 8, and 9 toward a conception of game logic as analyzing a Theory of Play rather than just games. We have shown how tools are available for such a program in our dynamic logics of information and preference that help paint a much richer picture of reasoning about social interaction. Still, we also considered objections to the resulting diversity, and problems with drawing natural boundaries. The resulting enterprise stands in need of philosophical reflection as much as technical development, but we hope to have shown the interest of both.

10.11 Literature

This chapter is based on van Benthem (2011b), van Benthem (2011d), van Benthem (2011f), and van Benthem et al. (2011).

Papers with a strong influence on the above views are Bicchieri (1988), Stalnaker (1999), Halpern (2001), and Halpern (2003a), as well as the game-theoretical literature mentioned at several places in this text: see Perea (2012) and Brandenburger et al. (2014) for congenial recent sources.