
3

Games with Imperfect Information

Games do not just consist of moves and outcomes that can be evaluated; they are also *played*, and focusing on players naturally involves their knowledge and information. In particular, in games of imperfect information such as card games, or many natural social scenarios, players need not know exactly where they are in the game tree. Reasons for this may be diverse: limited observation (say, hidden information in card games), processing limitations (say limited memory), or yet other factors. Thus, in addition to the actions and preferences studied so far, we need knowledge and eventually also belief, drawing once more on ideas from philosophical logic. In this chapter, we will show how games of imperfect information fit with standard epistemic logic, and then resume some of our earlier themes of Chapters 1 and 2 in this richer setting.

A conspicuous trend in all this is the emergence of players and play as an object of study in its own right. Already in Chapter 1, while game forms were just a static playground of all possible actions, they were navigated by actual players as they chose their strategies. Chapter 2 then added more information about how players viewed this playground in terms of their preferences. In this chapter, we emphasize one more aspect of players, namely their information processing as a game proceeds. All of this will eventually lead to the dynamic-epistemic logics of Part II providing a Theory of Play.

3.1 Varieties of knowledge in games

In this chapter, we will mainly use knowledge as modeled in epistemic logic (cf. Fagin et al. 1995, van Benthem 2010a). This is essentially the fundamental notion of an agent having semantic information that something is the case, referring to

the current range of possibilities for the actual situation.¹⁵ Foundations of game theory also involve the beliefs of players, and so do more general logics of agency, but doxastic logic will only be a side theme in this chapter, coming into play mostly in Part II of this book, by enriching the epistemic models presented here.

Perfect information Knowledge arises in games in different ways. Most scenarios discussed so far in this book are games of perfect information, where, intuitively, players know exactly which node of the game tree they are at as play proceeds. This corresponds to several informational assumptions: players know the game they are playing, and also, their powers of observation allow them to see exactly what is going on as play unfolds. But even then, in a branching tree, there is uncertainty as to the future: the players do not know which history will become actual as play proceeds. In this sense, the modal logic of branching actions in Chapter 1 is already an epistemic logic talking about possible future continuations of the current stage, and one can think of Backward Induction as a way of predicting the future in this setting of uncertainty. But for now, we discuss another important sense in which knowledge enters games.

Imperfect information In games of imperfect information, ignorance gets worse, and players need not know their exact position in a game tree as play proceeds. As before, this can be analyzed at two levels: that of local actions, and that of powers over outcomes, each with their own notions of structural invariance. We will focus on the former, deferring the latter view mainly to Chapter 11. Imperfect information games support an epistemic language, and with actions added, a combined modal-epistemic logic.

Varieties of knowledge in games Yet other kinds of knowledge are relevant to players, less tied to the structure of the game, such as knowledge about the strategies of others. We will discuss all varieties mentioned here: procedural, observational, and multi-agent, in Chapter 6, and once again in Chapter 10 when exploring our Theory of Play.

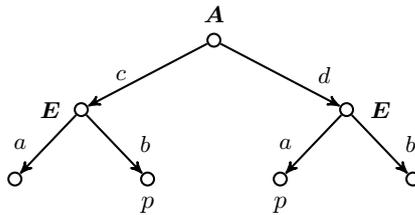
¹⁵ Other views of knowledge, based on fine-grained information closer to the syntax of propositions, are also relevant to understanding games, for instance, in studies of play that involve awareness (cf. Halpern & Rêgo 2006). However, they will remain marginal in this book, except for a brief appearance in Chapter 7.

3.2 Imperfect information games at a glance

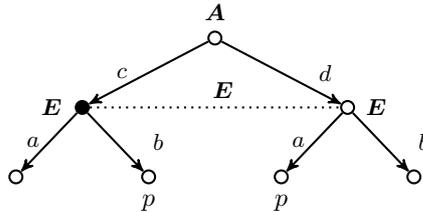
Games of imperfect information are process graphs as in Chapter 1, but with a new feature of uncertainty marking between certain nodes.

EXAMPLE 3.1 Game trees with and without uncertainty

First consider a typical picture of an extensive game as before:



Here, at the root, a modal formula $[c \cup d](a \cup p)$ expressed player E 's having a strategy for achieving outcomes satisfying p . To deal with imperfect information, game theorists draw dotted lines, whose equivalence classes form "information sets." Consider the above game with a new feature, an uncertainty for player E about the initial move played by A . Perhaps A put the initial move in an envelope, or E did not pay attention ...



Intuitively, E 's situation has changed considerably here. While the old strategy for achieving p still exists, and E knows this, it is unclear whether to choose *left* or *right*, since the exact location is unknown. E 's powers have diminished. ■

Books on game theory are replete with more sophisticated examples of imperfect information games (cf. Osborne & Rubinstein 1994), and further illustrations will occur in the course of this chapter, and at various places later on in this book.

3.3 Modal-epistemic logic

Process graphs with uncertainty Games of imperfect information are process graphs as in Chapter 1, but with a new structure among states, as brought out in the following notion.

DEFINITION 3.1 Epistemic process graphs

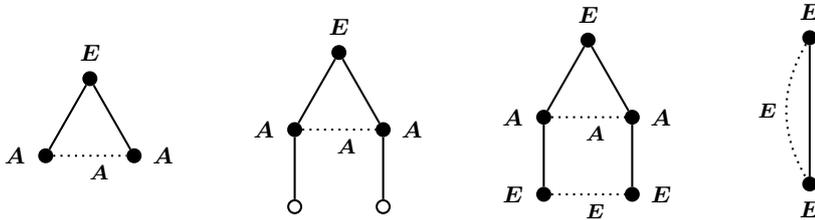
Epistemic process graphs are of the form $M = (S, \{R_a \mid a \in A\}, \{\sim_i \mid i \in I\}, V)$ consisting of game models as in Chapter 1, with added binary equivalence relations \sim_i for players i that represent their uncertainties as to which state is the actual one. When relevant, we also mark one particular state s in S as the actual one. ■

Thus, the relations \sim_i encode when player i cannot tell one node from the other over the course of the game. Epistemic process graphs can be generalized to arbitrary binary epistemic accessibility relations in a standard fashion, but we will have no need of this generality for the points to be made in this chapter. We can also add preference relations as in Chapter 2, where, of course, players may have quite different preferences between epistemically possible states (cf. Liu 2011).

In principle, any sort of uncertainty pattern might occur in epistemic process graphs. At various stages, players need not know what the opponent has played, what they have played themselves, whether it is their turn, or even whether the game has ended. This soon takes us beyond the usual game-theoretic setting where uncertainty links only run between players' own turns. (Uncertainty about what others will do is explored in Battigalli & Bonanno 1999a.)

EXAMPLE 3.2 Further imperfect information games

Plausible scenarios for the following pictures are left as an exercise for the readers.



We will take these diagrams in great generality, since epistemic logic over arbitrary structures fits well with the literature on imperfect information games (cf. Bonanno 1992a, Battigalli & Bonanno 1999a, and Battigalli & Bonanno 1999b). ■

REMARK General models

It is important to note that epistemic process graphs are just one way of associating epistemic models with games. As we will see in Chapters 5 and 6, and in Part II, it is equally feasible to take other objects as states, say complete histories of a game or more abstract possible worlds, to model other forms of information that players may have as they navigate a game. Indeed, general models $\mathbf{M} = (S, \{\sim_i \mid i \in I\}, V)$ of epistemic logic leave the choice of the set of states S completely free, subsuming all of these special cases. The epistemic definitions to follow in this chapter work in this full generality, even where stated only for epistemic process graphs.

Modal-epistemic language Epistemic process graphs are models for a joint modal-epistemic language. All the modal formulas of Chapter 1 still make sense for the underlying action structure, but now we can also make more subtle assertions about players’ adventures en route, using the dotted lines as a standard accessibility relation for epistemic knowledge operators interpreted in the usual way.

DEFINITION 3.2 Truth definition for knowledge

At stage s in a model, a player *knows* exactly those propositions φ that are true throughout the information set of s , i.e., at all points connected to s by an epistemic uncertainty link:

$$\mathbf{M}, s \models K_i \varphi \quad \text{iff} \quad \mathbf{M}, t \models \varphi \text{ for all } t \text{ with } s \sim_i t.$$

This says that formula φ is true all across agent i ’s current semantic range. ■

This epistemic logic view applies at once to imperfect information games.

EXAMPLE 3.3 Ignorance about the initial move

In the imperfect information game depicted earlier, after \mathbf{A} has played move c in the root, in the state marked by the black dot (in fact, in both states in the middle), \mathbf{E} knows that playing move a or b will give p , as the disjunction $\langle a \rangle p \vee \langle b \rangle p$ is true at both middle states. This may be expressed by the epistemic formula

$$K_{\mathbf{E}}(\langle a \rangle p \vee \langle b \rangle p)$$

On the other hand, there is no specific move of which \mathbf{E} knows that it guarantees a p -outcome, which shows in the black node in the truth of the formula

$$\neg K_{\mathbf{E}}\langle a \rangle p \wedge \neg K_{\mathbf{E}}\langle b \rangle p$$

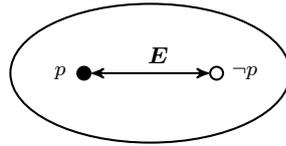
This is the famous “de dicto” versus “de re” distinction from philosophical logic.

Such finer distinctions are typical for a language with both actions and knowledge for agents. They also occur in fields like philosophy, AI, and computer science. We will see similar patterns in the epistemic temporal logics of Chapters 5 and 9.

Iterations and group knowledge An important and characteristic feature of epistemic languages is iteration. Players can have knowledge about each other’s knowledge and ignorance via formulas such as $K_{\mathbf{E}}K_{\mathbf{A}}\varphi$ or $K_{\mathbf{E}}\neg K_{\mathbf{A}}\varphi$, and this may be crucial to understanding a game. Indeed, very basic informational episodes such as asking a question and giving an answer crucially involve knowledge and ignorance about the information of others.

EXAMPLE 3.4 A model for a question/answer scenario

A question answer episode might start as follows. Agent \mathbf{E} does not know whether p is the case, but \mathbf{A} is fully informed about it ($\sim_{\mathbf{A}}$ is just the identity relation). The black dot indicates the actual world, the way things really are:



In the black dot to the left, the following epistemic formulas are true

$$p, K_{\mathbf{A}}p, \neg K_{\mathbf{E}}p \wedge \neg K_{\mathbf{E}}\neg p, K_{\mathbf{E}}(K_{\mathbf{A}}p \vee K_{\mathbf{A}}\neg p) \\ C_{\{\mathbf{E}, \mathbf{A}\}}(\neg K_{\mathbf{E}}p \wedge \neg K_{\mathbf{E}}\neg p), C_{\{\mathbf{E}, \mathbf{A}\}}(K_{\mathbf{E}}(K_{\mathbf{A}}p \vee K_{\mathbf{A}}\neg p))$$

Now \mathbf{E} can ask \mathbf{A} whether p is the case: \mathbf{E} knows that \mathbf{A} knows the answer. ■

The new notation $C_{\{\mathbf{E}, \mathbf{A}\}}$ is significant here. Communication and playing games are forms of shared agency that create groups of agents with knowledge of their own. A ubiquitous instance is the following notion.

DEFINITION 3.3 Common knowledge

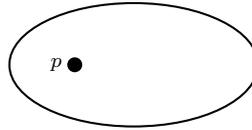
A group G has *common knowledge* of a proposition φ in an epistemic model \mathbf{M} at state s if φ holds throughout the reachable part of the relevant epistemic model.

Formally, $\mathcal{M}, s \models C_G \varphi$ iff φ is true in all those states that can be reached from s in a finite number of \sim_i steps, where the successive indices i can be any members of the group G . ■

In the above imperfect information game, \mathbf{E} 's plight was common knowledge between the players. Here is a scenario showing how common knowledge may arise:

EXAMPLE 3.4, CONTINUED Epistemic dynamics

There is a dynamics to the earlier question that can be modeled separately. Intuitively, the truthful answer “Yes” is an event that changes the initial model, taking it to the following one-point model where maximal information has been achieved:



The common knowledge formula $C_{\{\mathbf{E}, \mathbf{A}\}} p$ now holds at the black dot. ■

We will not make much use of iterations in this chapter, but they will occur at many places in this book, as they are crucial to understanding social interaction. The latter theme is of course much broader than what we have just shown (see van Ditmarsch et al. 2007, van Benthem 2011e for logical studies of communicative acts), and Chapter 7 will take up this theme in much greater generality.

Uniform strategies and nondeterminacy Let us now return to the topic of actions. A striking aspect of the preceding imperfect information game is its nondeterminacy. Uncertainty links have effects on playable strategies. Telling player \mathbf{E} to do the opposite of player \mathbf{A} was a strategy guaranteeing outcome p in the original game, but, although still present, it is unusable now. For, \mathbf{E} cannot tell what \mathbf{A} did. We can formalize this notion using a special kind of strategies.

DEFINITION 3.4 Uniform strategies

A strategy in an imperfect information game is *uniform* if it prescribes the same move for a player at epistemically indistinguishable turns for that player. ■

This restriction has an epistemic flavor, as players must know what the strategy tells them to do. In Example 3.1, neither player has a uniform winning strategy, interpreting p as the statement that player \mathbf{E} wins. Player \mathbf{A} did not have one to begin with, and \mathbf{E} has lost the one that used to work.

We will discuss the logical form of uniform strategies in more detail in Chapter 4, including the epistemic issue of what players know about their effects.

Calculus of reasoning In addition to defining basic notions, a logic of games should provide the means for analyzing systematic reasoning about players’ available actions, knowledge, and ignorance in arbitrary models of the above kind.

FACT 3.1 The complete set of axioms for validity in modal-epistemic logic is:

- (a) The minimal modal logic for each operator $[A]$.
- (b) Epistemic $S5$ for each knowledge operator K_i .

With a common knowledge operator added to the language, we get a minimal logic incorporating principles for program iteration from propositional dynamic logic (cf. Fagin et al. 1995).¹⁶ In Chapter 5, we will see it at work when analyzing the behavior of programs for uniform strategies.

3.4 Correspondence for logical axioms

Our minimal epistemic action logic is weak. In particular, it lacks striking axioms linking knowledge and action. Still, some commutation principles look attractive, as we shall see now.

EXAMPLE 3.5 Interchanging knowledge and action
Consider the modal-epistemic interchange law

$$K_i[a]\varphi \rightarrow [a]K_i\varphi$$

This seems valid for many scenarios. A person knows a priori that after dropping a full teacup, there will be tea on the floor. And after having dropped this cup, the person knows that there is tea on the floor. Even so, this was not in our minimal logic, since it can be refuted. A person may know that drinking too much makes people boring, but after drinking too much, the person is not aware of being boring.

¹⁶ One might extract a few further special-purpose axioms from Osborne & Rubinstein (1994). The fact of who is to move is taken to be common knowledge between players: $turn_i \rightarrow C_{\{1,2\}}turn_i$. Also, all nodes in the same information set have the same possible actions: $\langle a \rangle \top \rightarrow C_{\{1,2\}}\langle a \rangle \top$. Both properties will reappear in later chapters.

This actually tells us something interesting. The axiom is valid for an action without “epistemic side-effects,” but it may fail for actions with epistemic import. ■

This is not just an amusing example. In games, players often know at the start what particular strategies may achieve, and of course, they also want to have that knowledge available as they play successive moves of these strategies.

Similar observations hold for the converse $[a]K_i\varphi \rightarrow K_i[a]\varphi$. It fails in the imperfect information game of Example 3.1. In the black intermediate point, \mathbf{E} 's going right will reveal afterward that p holds, even though \mathbf{E} does not know right now that going right leads to a p -state.

But then, this game is a strange scenario, in that the uncertainty in the middle has suddenly evaporated at the end. Indeed, the law $[a]K_i\varphi \rightarrow K_i[a]\varphi$ is valid in games where agents learn only by observing new events and not suddenly by a miracle. To remove the miracle in the above scenario, one could add an explicit information-producing action at the end of seeing where we are, a theme that will return at the end of this chapter.

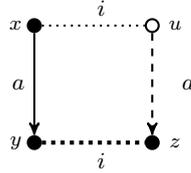
Correspondence analysis of special axioms We can be more precise here, using a well-known technique. What the above logical axioms say about the special games where they are valid can be stated exactly in terms of modal frame correspondences. Typically, what we find then are special assumptions about epistemic effects of actions in games, and players' abilities.

In what follows, we say that a formula holds in a graph if it is true at all points in that graph under all valuations for its proposition letters.

FACT 3.2 $K_i[a]\varphi \rightarrow [a]K_i\varphi$ holds in an epistemic process graph \mathbf{G} iff \mathbf{G} satisfies the following property: $\forall xyz : ((xR_a y \wedge y \sim_i z) \rightarrow \exists u : (x \sim_i u \wedge uR_a z))$.

Proof (a) If the property holds, then so does the axiom. For any valuation V on \mathbf{G} , suppose that $\mathbf{M} = (\mathbf{G}, V), s \models K_i[a]\varphi$. Now consider any point t with $sR_a t$: we show that $\mathbf{M}, t \models K_i\varphi$. So, let v be any world epistemically related to t (i.e., $t \sim_i v$). By the given property, there is also a point u with $s \sim_i u \wedge uR_a v$. But then, by the assumption about s , we have $\mathbf{M}, u \models [a]\varphi$, and hence $\mathbf{M}, v \models \varphi$. (b) Assume that the interchange axiom holds at point s in \mathbf{G} . Now, we choose a special valuation V making φ true only in those worlds w that satisfy the condition $\exists u : s \sim_i u \wedge uR_a w$. Clearly, we have $\mathbf{M} = (\mathbf{G}, V), s \models K_i[a]\varphi$. Thus, it follows from our assumption about the axiom that also $\mathbf{M}, s \models [a]K_i\varphi$. This means that, for any t and v with $sR_a t \wedge t \sim_i v$, $\mathbf{M}, v \models \varphi$. But by the definition of V , this says that $\exists u : s \sim_i u \wedge uR_a v$, proving the stated property. ■

More graphically, this relation between an action a and the information of an agent i expresses a well-known property of “confluence” for two binary relations:



The commuting diagram says that agent i ’s performing action a can never create new uncertainties for i : all dotted lines after a must have been inherited from before.

This property has a clear import for games. It says that the player has *perfect recall*, in the following strong sense: one remembers what one knew before. Indeed, new uncertainty can only arise when comparing two different events a and b , which can only happen if a player makes a partial observation of a current action, and does not know just which one. This could be because the source of the action is another player, as in our earlier running example, or for more general reasons of privacy in communication. We will analyze scenarios of the latter sort in much greater detail in Part II of this book.¹⁷

REMARK Learning by observation only

The converse axiom $[a]K_i\varphi \rightarrow K_i[a]\varphi$ can be analyzed in the same style, leading to a corresponding property $\forall xyz : ((x \sim_i y \wedge yR_az) \rightarrow \exists u : (xR_au \wedge u \sim_i z))$. Its matching commuting diagram says that old uncertainties are inherited under publicly observed actions a . A generalized form of this property is sometimes called “no miracles” (cf. Chapter 9): old uncertainties of an agent i only disappear when new events make distinctions.

General logical methods The above proof exhibits a simple pattern that can be generalized. Many similar equivalences follow by standard modal techniques (cf. Blackburn et al. 2001). Indeed, frame correspondences between modal-epistemic axioms and constraints on imperfect information games need not be found ad hoc.

¹⁷ There are also other versions of perfect recall in the game-theoretic literature: see Bonanno (2004a) and Bonanno (2004b) for some alternatives to the current version, expressible in epistemic-temporal logics.

They can be computed using effective algorithms, because many natural axioms concerning games have so-called “Sahlqvist forms.”

Coda: Diversity of players Highlighting the role of players as information-processing agents, as done in this chapter, opens up a natural dimension of diversity for players themselves. We have emphasized perfect recall, but many other kinds of players are possible, such as automata with limited memory: by now, these are perhaps the majority of the companions we interact with in our lives. The logics of this chapter are not all biased toward agents having flawless memory. Indeed, it is also easy to write axioms for memory-free agents that can only remember the last move that they observed.

FACT 3.3 Memory-free players satisfy $K\varphi \leftrightarrow \bigvee_e (\langle e^U \rangle \top \wedge U(\langle e^U \rangle \top \rightarrow \varphi))$.

Using a universal modality U , and a backward-looking converse modality $\langle e^U \rangle$ over past actions or events, this formula says that an agent knows that φ is the case if φ is true in all worlds of the model that arise from the same preceding action.¹⁸

3.5 Complexity of rich game logics

Logics of games may look simple, but they do combine many different notions, such as action, preference, and knowledge. Now it is known that logics with combinations of modalities can have surprising complexity for their sets of valid principles. We first raised this theme in Section 2.4 of Chapter 2, but now elaborate on this topic (see also Chapter 12). What determines this behavior is the mode of combination.¹⁹ In particular, what may look like natural commutation properties can actually drive up complexity immensely, away from decidable logics to undecidable ones.

THEOREM 3.1 The minimal logic of two modalities $[1]$, $[2]$ satisfying the axiom $[1][2]\varphi \rightarrow [2][1]\varphi$, and with an added universal modality U , is undecidable.

In other words, simple decidable component logics may combine into undecidable systems, or sometimes even unaxiomatizable or worse, depending on the mode of

¹⁸ Liu (2008) studies agent variety in terms of memory and powers of observation and inference, a perspective that will become important in the Theory of Play of Chapter 10.

¹⁹ This high complexity was first shown in Halpern & Vardi (1989) for epistemic-temporal logics of agents with perfect recall. The paper also cataloged complexity effects of other relevant epistemic properties for a wide array of logical languages.

combination. For instance, the preceding bimodal logic is Π_1^1 -complete (i.e., the validities are as complex as the full second-order theory of arithmetic).

The technical reason for this explosion is that satisfiability in such logics may encode complex geometrical “tiling problems” on the structure $\mathbb{N} \times \mathbb{N}$ (see Harel 1985, Marx 2006, and van Benthem 2010b for mathematical details). Such an encoding works if the models for the logic have a structure with grid cells that looks enough like $\mathbb{N} \times \mathbb{N}$. The earlier commuting diagram imposed by our notion of perfect recall is precisely of this kind, and it is easy to show that epistemic action logics over game trees in a language with a universal modality, or common knowledge plus a future modality, can be non-axiomatizable. Details and further examples can be found in van Benthem & Pacuit (2006).

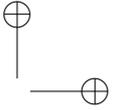
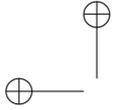
This complexity danger is sometimes made concrete in a simple picture. Modal logics of trees are harmless, while modal logics of grids are dangerous:



It is the looseness of trees that keeps their logics simple, and it is the close-knit structure of grids that makes logics complex. While extensive games are trees, adding further epistemic relations may induce grid structure, and the same may be true for adding preference order, as we saw in Chapter 2 with the complexity of rationality. Thus, while logic of games is a natural enterprise, it may actually lead to complex systems.

Discussion It is a moot point what these high complexity results mean. An optimistic line is that they are good news, since they show that logic of games is rich mathematically. Next, complexity need not always hurt: it may go down for important areas of special interest, such as finite games. Also, it has been suggested that the type of statement used in tiling reductions is not likely to occur in natural reasoning about interaction, so the part of the logic that is of actual use might escape the clutches of the complexity results.

Logic complexity versus task complexity But the discussion continues. Intuitively, there is a tension in high complexity results for logics of agents with perfect recall, since this property makes agents simple and well behaved. Likewise, commuting diagrams are devices that make reasoning smoother. While this is not a



paradox (a theory of regular objects may be richer than one for arbitrary objects), there is an issue whether complexity of a logical system tells us what we really want to know about the complexity of significant tasks faced by agents of the sort we are describing in that system.²⁰

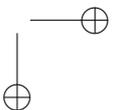
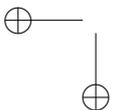
3.6 Uncovering the mechanics of the dotted lines

We conclude with a desideratum that will be taken up fully only later in this book. Our discussion of imperfect information games has served its purpose of linking up logics of games with epistemic logic, and bringing the players explicitly into the picture. But even so, it is not yet completely satisfactory. An imperfect information game gives us the traces of some *process* that left players with the uncertainties marked as dotted lines in the game tree. However, one is left guessing what that process may have been, and therefore, we have briefly discussed adding informational events of observation that remove “miracles.” Sometimes, it is hard to think of any convincing scenario at all that would leave the given traces. It would be far more attractive to analyze the sort of player and the sort of play that produces imperfect information games directly as a process in its own right. One expects a good deal of variety here, since information in games can come from so many different sources, as we have indicated at the beginning of this chapter.

Logical dynamics To do this, we need a shift in focus, to a logical dynamics (van Benthem 1996) making informational actions and other relevant events explicit. Games involve a wide variety of acts by players: prior deliberation steps, playing actual moves, observing moves played by others, perhaps acts of forgetting or remembering, and even acts of post-game analysis. Beyond these, the relevant dynamic repertoire may include acts of belief revision, preference change, or even of changing the current game. Making the logic of such events explicit is precisely what we will do in Part II of this book, using techniques from dynamic-epistemic logic (van Benthem 2011e).

Here is one instance of immediate relevance to this chapter. By analyzing the dynamics of information flow in games for concrete kinds of players, we see precisely what sorts of imperfect information games arise in natural scenarios. The

²⁰ In the setting of epistemic logic, this issue of internal task complexity versus external logical system complexity was taken up in Gierasimczuk & Szymanik (2011).



representation theorems to be proved in Chapter 9 provide concrete explanations for the characteristic traces left in a game by players with perfect recall endowed with powers of observation and policies of belief revision.

3.7 Conclusion

The main points Imperfect information games record scenarios where players have only limited knowledge of where they are in the game tree. We have shown how the resulting structures are models for a standard epistemic extension of the action logic of games that serves as their minimal logic. How agents traverse such games depends on special assumptions on their memory or powers of observation, and we have shown how perfect recall and other important properties of players (including bounded memory) are reflected precisely in natural additional modal-epistemic axioms via modal frame correspondences. We also noted that there can be hidden complexities in the system of validities, depending on the manner of putting component logics together.

Open problems A number of open problems come to light in our analysis. We already noted the need for uncovering the underlying procedural mechanics that produces the uncertainty lines in our game trees. We have also seen the need for a deeper analysis of the different kinds of knowledge that occur in games: either forward toward the future, or sideways or backward toward the past. Next, there is the obvious challenge of merging imperfect information and players’ preferences that drives the true dynamics of imperfect information games. Finally, our brief discussion of uniform strategies in imperfect information games highlighted the important distinction between “knowing that” and “knowing how”: a relatively neglected topic in logical studies of knowledge. Many of these topics will return in later chapters of Parts I and II of this book.

3.8 Literature

This chapter is based on van Benthem (2001b) and van Benthem & Liu (1994).

There is a vast literature on imperfect information in games and related social settings such as agency and planning, of which we mention Moore (1985), Fagin et al. (1995), Battigalli & Bonanno (1999a), van der Hoek & Wooldridge (2003), and Bolander & Andersen (2011).

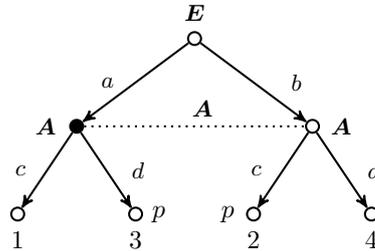
3.9 Further directions

As always, our chapter points at many further directions. We list a few.

Bisimulation and characteristic formulas Modal-epistemic logic is a game-internal language of action and knowledge. As in Chapter 1, then, one can line up two perspectives on its expressive power: (a) the roots of two games satisfy the same modal-epistemic formulas, (b) there is a bisimulation between the games linking the roots. Our earlier theory then generalizes. With imperfect information games, bisimulations need two back-and-forth conditions, one for actions and one for uncertainty links. Further, formulas of dynamic-epistemic logic characterize finite games up to bisimulation equivalence.

Uniform strategies We have seen how the notion of strategy changes in imperfect information games to having the same action at epistemically indistinguishable points. In Chapter 1, strategies were defined as programs in propositional dynamic logic. This approach is generalizable, but we now need the “knowledge programs” of Fagin et al. (1995), whose only test conditions for actions are knowledge statements. Knowing one’s strategy is related to the distinction between knowing that and knowing how. A strategy represents know-how for achieving certain goals, and this notion will be studied on its own in Chapter 4.

Powers and game equivalence As in earlier chapters, imperfect information games can be studied at different levels. In this chapter we have taken the fine-grained level of actions and local knowledge. One can also focus on strategic powers of players for influencing the outcomes of the game. To enter into the latter spirit, consider the following game. **E**’s two uniform strategies give powers $\{1, 3\}$ and $\{2, 4\}$, while **A**’s two uniform strategies LL and RR give powers $\{1, 2\}$ and $\{3, 4\}$:

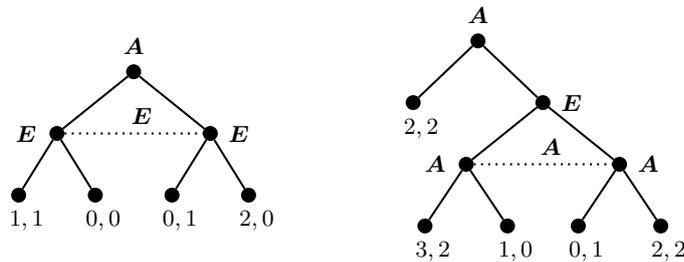


In terms of pure powers given by uniform strategies, this game tree is equivalent to one that interchanges the roles of the two players plus some outcomes. We refer to Chapter 21 for logical studies of imperfect information games like this, and once more to Thompson (1952), Kohlberg & Mertens (1986), and Bonanno (1992b) for game-theoretic studies of such settings. Within our logical framework, we will pursue this more global strategic level of game analysis in Chapter 11.

Fixed point logics for game solution In Chapter 2, the mechanics of game solution came out particularly clearly in fixed point logics. We have not given similar systems here, and indeed the fixed point theory of imperfect information games seems virtually nonexistent. One reason is the proliferation of equilibrium concepts in this area, and hence less of a grip on the logical essentials than what we found with Backward Induction.

Adding preferences An important feature left out altogether in this chapter is players’ preferences. While it is not hard to combine the logics of action and knowledge in the above with the logics of preference of Chapter 2, there are many subtleties in trying to understand the precise reasoning that is needed in concrete instances. Just to illustrate the entanglement of imperfect information and preference, recall a distinction made at the start of this chapter. We had knowledge linked to observation in imperfect information games, but also knowledge linked to uncertainty about how any game will proceed. We did not bring the two views together, but both arise in concrete settings.

Here are two scenarios, with our usual outcome order (*A*-value, *E*-value).²¹



²¹ The tree to the right is slightly adapted from an example in an invited lecture by Robert Stalnaker at the *Gloriclass Farewell Event*, ILLC Amsterdam, January 2010.

The game on the left suggests an extended use of Backward Induction, but the one on the right raises tricky issues of what **A** is telling **E** by moving right. We leave the question of what should happen in both games to the reader. Dégremon (2010) and Zvesper (2010) have more extensive discussion of what logic might say here.

Trading kinds of knowledge in games The issue of relating the different kinds of knowledge in games at the beginning of this chapter has not been settled. Many ideas are circulating about trading one kind for another: say, replacing uncertainty about the future by imperfect information-style uncertainty right now between different histories.²² We will return to combining and transforming notions of knowledge in Chapter 6, although eventually, we may just have to accept the diversity. Also, we may need other forms of knowledge beyond the purely semantic, such as syntactic awareness-based views (Fagin et al. 1995, van Benthem & Martínez 2008).

Adding beliefs As we have noted several times, games are not just driven by knowledge, but also by beliefs. A common representation for this richer setting uses epistemic models, adding a relation of comparative plausibility on epistemically accessible worlds (cf. Chapter 7). Belief is then what is true in the most plausible accessible worlds. This adds fine structure to all earlier topics, but we leave its exploration to Part II of this book.

Adding probability More expressive quantitative logics for numerical degrees of belief or frequency information about process behavior arise when we add probability functions to our models. Logic and probability mix well in the study of agency, as can be seen in the systems of Halpern (2003b) and van Benthem et al. (2009b), but in this book, we will ignore probabilistic perspectives, regarding them as an extra for later.

Game logics in other areas While the logics in this book are tailored to games, their scope is more general. Many points in the preceding apply to a much broader class of social scenarios. In particular, our modal-epistemic logic also occurs in planning, when agents act under uncertainty (Moore 1985). There are also analogies between what we do in this book and recent work on dynamic-epistemic logics for multi-agent planning (cf. Bolander & Andersen 2011, Andersen et al. 2012).

²² Compare this to the Harsanyi Doctrine in game theory (Osborne & Rubinstein 1994).