# Games in Dynamic-Epistemic Logic

Johan van Benthem
University of Amsterdam & Stanford University

in G. Bonanno & W. van der Hoek, eds., *Bulletin of Economic Research* 53:4, 219–248 (Proceedings LOFT-4, Torino).

## Abstract

We discuss games of both perfect and imperfect information at two levels of structural detail: players' local actions, and their global powers for determining outcomes of the game. We propose matching logical languages for both. In particular, at the 'action level', imperfect information games naturally model a combined 'dynamic-epistemic language' – and we find correspondences between special axioms and particular modes of playing games with their information dynamics. At the 'outcome level', we present suitable notions of game equivalence, plus some simple representation results.

## 1 Games as structures for logical languages: a first glance

**1.1 Modal logic of actions** The pictures in a typical game theory text have an immediate appeal to a logician. In particular, an extensive game tree depicts a many-agent process, where nodes are possible stages, with available actions for players indicated, and perhaps also with markings for relevant properties of these stages. This is exactly like *process graphs* ('Kripke models') for *modal logic* or related languages that are used by logicians and computer scientists for describing action structures. E.g., consider the following game tree for two players **A**, **E** with four possible actions c, d, a, b, and some special property p holding at two of the four possible end states:

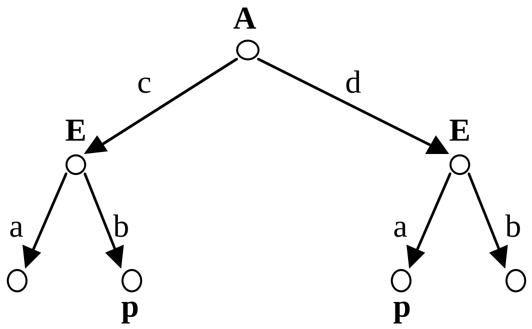

Fig. 1

Here is a typical modal assertion which is true at the root:

| | |
|---|---|
| $[c \cup d]\langle a \cup b\rangle p$ | each of the actions c and d leads to a state where either a or b can be executed to get to a final state where p holds. |

This modal formula says in the given process graph that player **E** has a *strategy* ensuring that the outcome of the game satisfies p . E.g., p might be the property that **E** wins, in which case the above modal formula expresses that ' **E** has a winning strategy'. More complex strategies than this 'single response' can be analysed with longer $[]\langle\rangle$ sequences of modal operators. This strategic feature is one of the things which makes extensive games of such great interest to logicians and computer scientists, as these model processes, not with dumb machines, but between intelligent agents and their complex interactions over time.

Now, the above way of thinking about game trees may not be *quite* the interpretation that game theorists themselves have in mind. The aim of this paper is to develop the logical view in a bit more detail, hoping that there is enough difference to make for new insights, but also enough resemblance to allow for meaningful communication.

**1.2 Imperfect information: adding epistemic logic** The modal view is only a first step, as it merely deals with pure game forms. Other features of real games suggest richer logical languages. Our main topic in this paper is *imperfect information*, where players may be uncertain in which state they are in the course of a game. To model this, game theorists draw 'dotted lines', or 'information sets'. Consider again the above game, but now with an uncertainty for player **E** about which move was played by **A** initially. (Perhaps, **A** put his move in an envelope.) To a logician, this is a typical model for a combined *modal-epistemic* language:

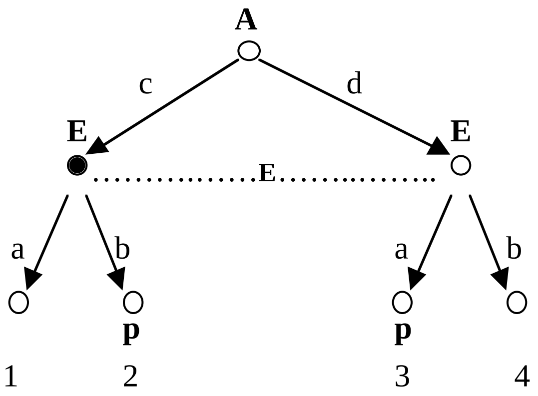

Fig. 2

The earlier modal formula $[c \cup d]\langle a \cup b\rangle p$ is still true at the root – since the action pattern is the same. But this time, we can make more subtle assertions about players' adventures 'en route', using the dotted lines as an accessibility relation for epistemic *knowledge operators* in the usual way. At any stage s , a player 'knows' exactly those propositions that are true throughout the information set to which s belongs. This epistemic connection has also been known to game theorists, ever since the 1970s. E.g., after **A** plays move c in the root, in the resulting 'black' state, **E** 'knows' that playing either a or b will give her p – since the disjunction $\langle a\rangle p \vee \langle b\rangle p$ is true at both 'middle states'. This may be expressed by means of the epistemic formula

$$K_E(\langle a\rangle p \vee \langle b\rangle p)$$

On the other hand, there is no *specific* move of which **E** knows that it guarantees an outcome satisfying p – which shows in the black node in the truth of the formula

$$\neg K_E\langle a\rangle p \;\&\; \neg K_E\langle b\rangle p$$

Think of the tragic someone who knows that the right partner is walking around right in this city, but does not know of any particular person whether (s)he is that partner.

Such finer distinctions are typical for a language describing both actions and knowledge for agents. They also occur in other fields: e.g. with planning in AI – so there is nothing exotic to our way of reading the above diagrams. These generalize process graphs to more sophisticated computational systems, like those on the Internet to-day, where agents in a group have the indicated actions available in principle at each stage of the process, but must operate under epistemic uncertainties.

**1.3 Preferences** Realistic games have still further essential structure, in particular, *preferences* of players between the different possible outcomes. These give rise to a next level of logical formalization, say by using the preference relations to interpret 'preference modalities'. But this evaluational structure will not be addressed in this paper, though it is certainly compatible with all that will be proposed below.

**1.4 Game equivalence: actions or outcomes?** What are key issues in a logical perspective on games? In this paper, we concentrate on two basic themes. The first is this. Choosing a language with a certain expressive power for describing the internal structure of games, or any class of models for that matter, is just one side of a coin. The other side is choosing a matching notion of *structural equivalence* between different presentations of a game. For instance, compare the following two games:

Fig. 3

These are not the same, intuitively, in their extensive structure of possible *actions*. E.g., on the left, the game can reach a state where **E** can choose between achieving outcome 2 or 3. No such stage occurs in the game on the right. Thus, at the level of possible *actions*, these games are intuitively different.

But when we merely look at possible *outcomes*, this verdict changes. In the left-hand game, player **A** has two strategies "left" and "right", which guarantee outcomes for the game in the two sets {1}, {2, 3} , respectively. Likewise, **E** has two strategies, which guarantee outcomes in {1, 2} and {1, 3} . We think of these sets as 'powers' that players have for forcing specific outcomes of the game. A set with more elements says the power is not strong enough for a unique outcome, just some 'upper range':

| | |
|---|---|
| **A**'s powers | {1}, {2, 3} |
| **E**'s powers | {1, 2}, {1, 3} |

Now, when we compute powers in the right-hand game, we obtain the same outcomes! First **E** has two strategies, which force {1, 2} and {1, 3} . Next **A** has 4 strategies LL, LR, RL, RR which yield, respectively, {1}, {1, 3}, {2, 1}, {2, 3} . Of these four, {1, 3} and {2, 1} can be dropped, as they represent weaker powers than {1}, and hence are redundant. Thus **A** has the powers {1}, {2, 3}, just as in the left-hand game.

Thus, game equivalence depends on how much structure one wants to describe: with a spectrum running from 'just outcomes' to 'all actions'. This is a virtue, not a vice. Process logics in computer science, or grammars in linguistics have the same options. E.g., a process may be described as a 'black-box' merely in terms of input-output outcomes, or in terms of its internal choices and other actions. In specific applications, the 'level of identification' will depend on the practical purpose at hand. The same seems true for games. In what follows, we mainly describe games at their local 'action level', but we will make excursions towards the global 'outcome level' as well.

Game equivalences also make sense with imperfect information. In terms of 'actions', one can now compare both physical game moves and epistemic 'uncertainty steps' across two games, looking for enough similarity in both. As for 'outcomes', we shall see later that, in terms of the appropriate powers for players in that setting, the game in Figure 2 is outcome-equivalent to the game tree in Figure 4 below, which has the roles of the two players plus some outcomes suitably interchanged:

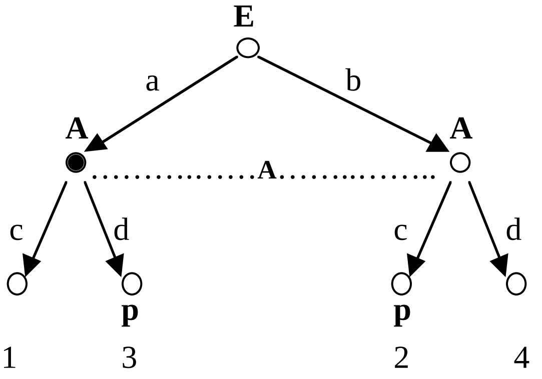

Fig. 4

**1.5 General logics and special axioms** The second main logical theme is this. Once we have chosen a particular language, at any ‘level of detail’, we get a set of universally *valid formulas*, those that are true at every state in any model whatsoever. The resulting minimal logics are well-known for modal and epistemic models, being respectively the poly-modal versions of ‘K’ and ‘S5’. But on top of such base systems, one finds special logics that are theories of special model classes satisfying restrictions. For instance, if we demand that players’ moves a are *partial functions*, then the modal logic acquires an extra axiom

$$\langle a\rangle\phi \rightarrow [a]\phi$$

We will encounter more such special axioms for the interplay of modal and epistemic operators in our study of imperfect information games, such as the interchange law

$$K_i[a]\phi \rightarrow [a]K_i\phi$$

(cf. Section 3.3). Such principles reflect special assumptions about players’ abilities, or the uncertainty inherent in the game. They yield different ‘logics’ for reasoning about the behaviour of players in different types of imperfect information game.

**1.6 Overview of the paper** Section 2 introduces some basic logical notions for perfect information games, viewed as process graphs, at both action and outcome levels. We use these in Section 3 to deal with actions and knowledge in imperfect information games, viewed as many-agent processes with uncertainty. The results are used in Section 4 to analyze several styles of playing such games, and in Section 5, to analyze outcome equivalence. Section 6 is a discussion of ways to go from here.

## 2 Extensive games as models for dynamic logics

We proceed in two stages, starting our logical analysis with *finite perfect information games*. This section is a brief survey of some basic logical notions to be used later on.

**2.1 Action languages** As we noted, perfect information game trees are models

$\mathbf{M} = (S, \{R_a \mid a \in A\}, V)$

where

(a) S is a set of *states*, (b) the $R_a$ are binary relations encoding the possible *transitions* for action/move types a , and (c) V is a *valuation* function for proposition letters denoting local properties of states.

In this most general logical set-up, games may have some distinguished vocabulary. In particular, propositional atoms $turn_i$ and $win_i$ say that player **i** is to move or wins, *end* holds at just the end-points, while other atomic statements may indicate values of final states to players. These models are described by a *dynamic modal language* with basic actions a, b, … and operations of

| | | |
|---|---|---|
| (a) | union $\cup$ | choice |
| (b) | composition **;** | sequential execution |
| (c) | Kleene star $^*$ | arbitrary finite iteration |
| (d) | $(\phi)\mathbf{?}$ | test if assertion $\phi$ holds |

Thus we can define further transition relations in a game, such as the single choice $a \cup b$, or the iterated choice $(a \cup b)^*$ which allows us to walk from a state to any state reachable from it by some finite sequence of a and b moves. For all resulting actions A (basic or compound) the dynamic language has modalities

$\langle A\rangle\phi$ "in some A-successor, $\phi$ holds"
$[A]\phi$ "in all A-successors, $\phi$ holds"

As we have seen above, this allows us to express single-response strategies, such as

$[c \cup d]\langle a \cup b\rangle p$

But one can also define much more complex assertions about possible actions and outcomes, using patterns of statement $[A]\langle B\rangle[C]\langle D\rangle \phi$ , etc. Here is another example of this expressive power. Consider the simplest use of *Backward Induction*, computing which player has the winning strategy at any node in a finite 2-player zero-sum game with outcomes marked $win_i$ or $\neg win_i$ for each player **i** . One starts in the end nodes, and then works upwards. Let A be the union of all available actions. Then the rule for the winning positions may be defined by the following recursion:

$$\mathrm{WIN}_i \leftrightarrow (end \,\&\, win_i) \vee (turn_i \,\&\, \langle A\rangle \mathrm{WIN}_i) \vee (\neg turn_i \,\&\, [A]\mathrm{WIN}_i)$$

This language can also speak about *strategies* explicitly. A strategy for player **i** is just a partial function from those nodes where it is **i**'s turn to concrete moves. Thus, a strategy $\sigma$ is a set of atomic transitions, just like any basic action. Then, for instance, saying that $\sigma$ is a winning strategy for **i** amounts to stating that

$$[A^*]\,(\mathrm{turn}_i \rightarrow \langle\sigma\rangle \mathrm{WIN}_i) \qquad \text{where } A^* \text{ goes from the root to any game node}$$

Using two strategy symbols in the same way, the language can analyze Nash equilibria (De Bruin 2000). Dynamic logic is decidable, both in general and over the special class of finite trees, and in both cases, there is a complete axiomatization for its set of valid principles. We refer to Blackburn, de Rijke & Venema 2000 for details.

**2.2 Bisimulation** The structural counterpart to the dynamic description language is the following notion of comparison between different games. It works on game trees, but equally well on any graph representation of possible actions and states:

Definition A *bisimulation* is any binary relation E between states of two graphs **M**, **N** such that, if x E y , then we have (1) atomic harmony, plus (2) zigzag clauses:

(1) x, y verify the same proposition letters
(2a) if x R z, then there exists u in **N** s.t. y R u and z E u
(2b) vice versa.

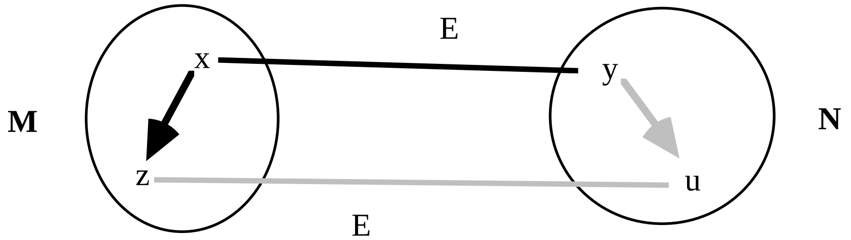

Fig. 5

Bisimulation is the typical 'action-similarity' for processes in logic and computer science. It preserves modal and dynamic formulas, and there are some converses, too. Here are two sample results that establish the link (cf. van Benthem 1996, 1998).

Theorem For finite graphs **M**, **N** with nodes s, t , the following are equivalent:

(a) **M**, s and **N**, t satisfy the same modal formulas

(b) there is a bisimulation E between **M**, **N** with sEt

This says that the language and the similarity relation match. The next result says that at this level of description, the language provides complete descriptions for any graph.

Theorem For each finite graph **M** and state s , there exists a dynamic logic formula $\delta(\mathbf{M}, s)$ such that the following are equivalent for all graphs **N**, t :

(a) $\mathbf{N}, t \models \delta(\mathbf{M}, s)$

(b) **N**, t is bisimilar to **M**, s

Similar definability and expressiveness results can be proved for all notions of game equivalence considered here (Barwise & Moss 1997, van Benthem 2000B). Thus, game structure can be described either via semantic equivalences between representations, or in terms of 'game formulas' in suitable formal languages (cf. Bonanno 1993). In the rest of the paper, we mainly concentrate on the latter method of description.

**2.3 Outcomes and powers** Moving now to outcomes, let us first sharpen up Section 1.4. *Strategies* in extensive games are defined as usual, being functions from all nodes where it is the relevant player's turn to actions that are available at that node. *Forcing relations* for a game then encode what players can achieve:

$\rho_G^{\mathbf{i}}$ sX *player* ***i*** *has a strategy for playing (the remainder of) game G from node s whose resulting states are always in the set of outcomes X*

Given this explanation, it is clear that forcing relations are *closed under supersets*:

(C1) *if* $\rho_G^{i}\, s, Y$ *and* $Y'$ *contains* $Y$, *then* $\rho_G^{i}\, s, Y'$

Another obvious constraint on these relations is *consistency*: players cannot force the game into disjoint sets of outcomes, or a contradiction would result:

(C2) *if* $\rho_G^{1}\, s, Y$ *and* $\rho_G^{2}\, s, Z$, *then* $Y, Z$ *overlap*

Moreover, all finite 2-player games are *determined*: for any winning convention, one of the two players must have a winning strategy. Stated in the present terminology, this is a condition of *completeness*. Let S be the total set of outcome states:

(C3) *if not* $\rho_G^{1}\, s, Y$ , *then* $\rho_G^{2}\, s, S–Y$; *and the same for* ***2*** *vis-á-vis* ***1***

Any finite perfect information game G between players **1**, **2** yields powers in the root s satisfying (C1), (C2), (C3). Conversely, these conditions are also *all* that must hold:

Proposition Any two families $F_1$, $F_2$ of subsets of some set S satisfying conditions (C1), (C2), (C3) are the root powers in some two-step game.

Proof Start with player **1** and let her choose between successors corresponding to the *inclusion-minimal* sets in $F_1$ . At these nodes, player **2** gets to move, and can choose any member of that set. Clearly then, player **1** has the powers specified in $F_1$ . In this game, player **2** can force any set of outcomes that overlaps with each of the sets in $F_1$ . But by the constraints, it is easy to see that these are precisely the sets in $F_2$ . For instance, if some set of outcomes A overlaps with all these sets, the complement –A cannot be among them, and hence A was in $F_2$ , by determinacy. **QED**

This result gives an outcome-level *normal form* for games, closely related to the usual 'strategic form'. It is also the 'distributive normal form' of standard propositional logic. Indeed, the usual Boolean operations that produce such normal forms form a *logical calculus of game equivalence* (van Benthem 1999). Figure 3 is a typical case: the game equivalence stated there was nothing but the law of Boolean Distribution:

$$p \mathbin{\&} (q \vee r) \leftrightarrow (p \mathbin{\&} q) \vee (p \mathbin{\&} r)$$

**2.4 A forcing language and its simulation** To describe games at this outcome level, one can introduce another logical language, whose key operator is this:

> $\{G, \mathbf{i}\}\phi$ is true at game node s if player **i** has a strategy for playing game G from there on which forces a set of outcomes all of which satisfy $\phi$

This is the 'game-internal' version of the game modalities in Parikh 1985, Pauly 2000. The corresponding notion of bisimulation between different games is as follows.

Definition A *power bisimulation* between two games G, G' is a binary relation E between game states in G, G' satisfying

(1) if x E y, then x, y satisfy the same proposition letters.

(2) if x E y and $\rho_G^{\mathbf{i}}$ x, U, then there exists a set V with $\rho_{G'}^{\mathbf{i}}$ y, V and $\forall v \in V \, \exists u \in U \; u E v$; and vice versa.

This can be extended naturally to forcing sets of intermediate positions in a game.

**2.5 Coda: internal versus external languages** Logical description of games, as that of processes, can work at various levels. The formal languages in this paper are *game-internal*. They define statements about stages of a play of a game. But the study of game equivalence also suggests a *game-external* point of view, with expressions for whole games. This calls for external languages, whose expressions denote games, using *game-forming operations* such as 'choice', 'role switch', or 'composition'. Distribution in Section 2.3 was an example: the expression $p\&(q\vee r)$ stood for a game of a certain shape, not for an assertion. The external viewpoint has its attractions too, e.g., when describing game equivalences in an algebraic manner. Moreover, the internal and external viewpoints can be combined – but in this paper, we will stick with the former.

## 3 Imperfect information games and dynamic-epistemic logic

In games of imperfect information, during play, players may not know exactly where they are in the tree. This may come about for different reasons. Players may have cognitive limitations, such as bounded memory or limited perception of others' moves. But there are also games that impose ignorance by their definition, such as many card or parlour games. Whatever the sources, one standard logical approach suggests itself.

**3.1 The epistemic language** Games of imperfect information add epistemic structure among states, in the form of binary '*uncertainty*' relations $\sim_i$ for players **i** . The resulting models contain a 'multi–S5' structure:

$$\mathbf{M} = (S, \{R_a \mid a\in A\}, \{\sim_i \mid i\in I\}, V)$$

The *equivalence relations* $\sim_i$ encode that player **i** cannot tell one node from the other when she gets to them in the course of the game. This extends the game-theorist's usual 'information sets', which only encode informational constraints for a player on those nodes where it is his/her turn. In principle, in these models **M** any uncertainty pattern might occur. Players need not know what the opponent has played, or what they played themselves, they need not know if it is their turn, or whether the game has ended, etc. One can think up plausible scenarios with all of the following pictures:

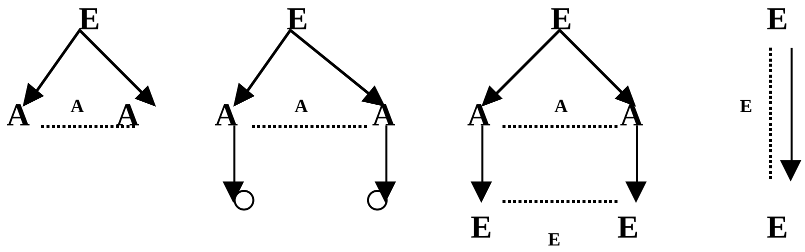

Fig. 6

Pictures like this are not particular to games. They also occur in accounts of planning, where agents must act in partial ignorance of the true state of affairs (Moore 1985). In what follows, we consider these game diagrams in great generality. Our aim in doing this is not to ‘improve’ on the rich existing game-theoretic literature on imperfect information games (cf. Osborne & Rubinstein 1994, Bonanno 1992, Battigalli & Bonanno 1999), or take a stand on any current debate. We merely wish to suggest a systematic logical perspective that may provide an interesting ‘second opinion’.

**3.2 Epistemic-dynamic logic** The above richer models naturally interpret a richer epistemic-dynamic language, which adds knowledge operators for each player:

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

Now we can talk about knowledge and ignorance of players when the game has reached a certain state. This was demonstrated in Section 1 using formulas like

$K_E(\langle a\rangle p \vee \langle b\rangle p)$
$\neg K_E\langle a\rangle p \ \& \ \neg K_E\langle b\rangle p$

referring to the black state in Figure 2. Here player **E** has a strategy for achieving p in principle, but it is useless in practice, as its advice (“play a if c was played by **A**, otherwise, play b”) depends on a distinction which **E** cannot resolve. In imperfect information games, players can only deliberately use

*uniform strategies*, whose next action at any node
is the same across their information set at that node

E.g., in the above picture, **E** has only two uniform strategies LL and RR, of which she does not know if they guarantee outcome p . (Nota bene. Game theorists *define* ‘strategies’ in imperfect information games as what we call uniform strategies here – but our distinction makes sense in logic, being similar to that between ‘constructive’ and more general ‘non-constructive’ proof.) Our example suggests that uniform strategies are those for wich players *know that they will work*. We will look at this later, because matters are more complicated than might seem from the above example.

A further important feature of epistemic languages are *iterations*. Players can have knowledge about each others' knowledge and ignorance (via formulas like $K_iK_j\phi$, $K_i\neg K_j\phi$ ), and this may be crucial to understanding the course of a game. Also, players may achieve *common knowledge* about certain facts, written as follows:

$C_{\{1,2\}}\phi$ *$\phi$ is true in all those states that can be reached from the current one in a finite number of $\sim_1$ and $\sim_2$ steps*

For instance, in the above game, **E**'s plight is common knowledge between the players. As for systematic reasoning about players' available actions, knowledge, and ignorance in arbitrary models of this kind, the complete set of *axioms for validity* in dynamic-epistemic logic is the following.

(a) the *minimal dynamic logic* for the modal operators [A]
(b) *epistemic S5* for each knowledge operator $K_i$

With a common knowledge operator added, we also get the minimal logic of that (cf. Fagin et al. 1995). There are no further axioms in general epistemic-dynamic logic.

**3.3 Constraints for special axioms** But there is more to be said. Specific imperfect information games may validate additional epistemic-dynamic principles. Osborne & Rubinstein 1994, Chapter 11, gives a number of such constraints. These induce special axioms for reasoning about players in games following a certain style.

Example 'The fact who is to move is common knowledge between players.' This is a reasonable requirement, e.g., in many parlour games. The relevant formula is

$$\text{turn}_i \rightarrow C_{\{1,2\}}\text{turn}_i$$

Example 'All nodes in the same information set have the same possible actions'. The corresponding dynamic-epistemic formula reads (with 'T' for *true*):

$$\text{turn}_i \;\&\; \langle a\rangle T \rightarrow K_i\langle a\rangle T \quad \text{or perhaps even} \quad \langle a\rangle T \rightarrow C_{\{1,2\}}\langle a\rangle T$$

As a final example, consider interactions between players' actions and their knowledge. *In general*, dynamic-epistemic logic validates no such principles.

Example The Interchange Principle $K_i[a]\phi \rightarrow [a]K_i\phi$

This says that, if player **i** knows right now that doing a will bring about $\phi$, then doing a will result in her knowing that $\phi$. This implication is reasonable for actions *without epistemic side-effects*, but not in general. It fails if action a is 'have one more beer', and $\phi$ the assertion that **i** is making a fool of himself. On game trees, the formula only holds under the following constraint on action a and **i**'s uncertainties:

$$\forall xyz: ((\, x\, R_a\, y \;\&\; y \sim_i z) \rightarrow \exists u: ((\, x \sim_i u \;\&\; u\, R_a\, z\,)$$

This is a well-known *confluence property* for the relations of action and uncertainty:

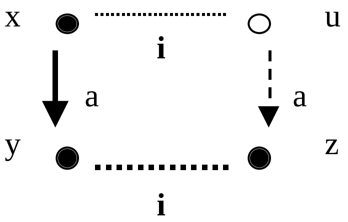

Fig. 7

The converse axiom $[a]K_i\phi \rightarrow K_i[a]\phi$ also fails in general. Reading this paper will make you know certain things $\phi$, even though you might not have known beforehand that it would do so. Both kinds of interchange principle will return in Section 4.3.

As a final technical point, such *correspondences* between dynamic-epistemic axioms and nice constraints on imperfect information games do not have to be found ad-hoc. They can be *computed* using well-known algorithms from modal logic (van Benthem 1996, Blackburn, de Rijke & Venema 2000).

### 3.4 Bisimulation and characteristic formulas

Dynamic-epistemic logic provides a game-internal language speaking about possible actions and uncertainties. Alternatively, as in Section 2.2, one can line up the following two perspectives:

(a) *the roots of two games satisfy the same dynamic-epistemic formulas*
(b) *there exists a* bisimulation *between the games connecting the roots*

In this setting, bisimulations will now have zigzag conditions, both for possible actions and for 'uncertainty jumps' inside information sets. Likewise, we can use formulas of dynamic-epistemic logic to characterize finite games up to bisimulation equivalence.

## 4 Ways of Playing Games

### 4.1 Statics versus dynamics

Game trees are static pictures of all that can happen in a game. But in any specific run, there is an actual course of events, taking players to successive game states, each with different knowledge and ignorance:

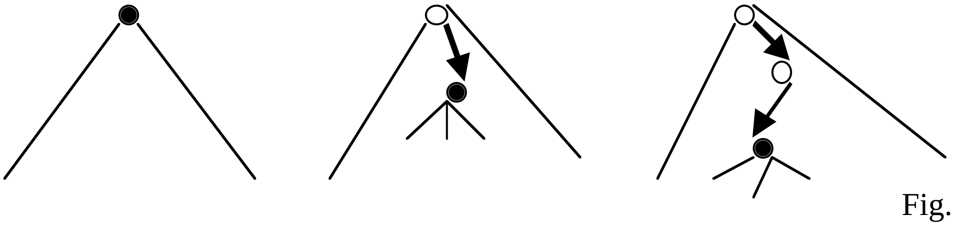

Fig. 8

This suggests a closer look at the logical 'dynamics' of game playing, in particular, the *update mechanisms* that produce or remove ignorance. Our dynamic-epistemic languages describe what is true at successive nodes, keeping track of the changes in players' information. But one can be more systematic. We can also think about *different modes of playing a game*, and determine their modal-epistemic properties.

**4.2 Perfect recall** An important mode of game playing satisfies *Perfect Recall*:

> Players remember their own previous moves,
> plus the uncertainties they had at each stage.

This notion is defined for a player **i** in Osborne & Rubinstein 1994 by the following demands on any two nodes related by its uncertainty relation $\sim_i$ :

> *going back along the unique histories determined by these nodes,*
> *one must find the same record of moves for player* ***i*** *when it is her turn,*
> *and the same sets of* ***i****– indistinguishable alternatives at each node.*

Analyzing this description stepwise, we find two semantic aspects to what happens.

(**a**) When **i** is to move at x , and plays a to arrive at game state y , then any $\sim_i$ – uncertainty z after that must arise as follows: " z was also reached via a , from some node u , and moreover, u and x must share the same uncertainty ". But the latter means no more or less than: $x \sim_i u$ . This may be pictured as a *commuting diagram*:

$$\begin{tikzcd}
\mathbf{i} & x \arrow[r, dash, dashed, "\mathbf{i}"'] \arrow[d, "a"] & u \arrow[d, dashed, "a"'] \\
 & y \arrow[r, dash, dotted, "\mathbf{i}"'] & z
\end{tikzcd}$$

Fig. 9

In terms of Sections 2.2, 3.4, a modal logician might restate this as follows:

(i) $\sim_i$ is a bisimulation: w.r.t *converse* actions $a^{\smile}$

(ii) **a** is a bisimulation: w.r.t. 'uncertainty jumps'.

or in the physicists' parlance, in this case, *action and uncertainty commute*

(**b**) Next, when **i** is not to move at x , there is still such a bisimulation, but now with the difference that we do not require the action to be the same on both sides. The reason is this. The only source of (increased) uncertainty that we acknowledge is lack of information *about another player's move*. In particular, the latter may happen starting from one single node, say the root, and the other player's move generates uncertainty straightaway. The diagram for this case is essentially like that of Figure 9.

Proposition Perfect Recall is satisfied precisely by those dynamic-epistemic models that satisfy the given two commutative diagrams for each player.

Proof *From PR to ‘commutation’*. Here is one case. Suppose **1** was to move at $x$, and $x\ R_a\ y \sim_1 z$. Because of PR for **1**, at $z$, the same last action $a$ must have taken place there, while at the beginning of the latter, say $u$, the same uncertainty set existed as at $x$. But the latter is equivalent to just: $x \sim_1 u$. The other case is similar. Here, players can recognize moves as being their own, or by the other player, and also, they know the nuber of moves plus turn-taking record along any history.

*From ‘commutation’ to PR*. This uses induction down any two paths of equal length down the tree whose final nodes are $\sim_I$-related. It is easy to see that these paths shared a point somewhere, while afterwards, the bisimulation diagrams for both players ensure that (a) they have the same moves for both players, (b) there were uncertainty links between corresponding stages from then on until the end. **QED**

The preceding Proposition provides a dynamic-epistemic *axiomatization* for imperfect information games with Perfect Recall (cf. Section 3.3.) In reasoning about players' adventures in games of this sort, we can use the earlier minimal dynamic-epistemic logic, but now enriched with the following interchange axioms:

(1) $\text{turn}_i \ \&\ K_i[a]\phi \rightarrow [a]K_i\phi$
(2) $\neg\text{turn}_i \ \&\ K_i[A]\phi \rightarrow [A]K_i\phi$
with $A$ the union of all actions available to the other player

Proposition The following two statements are equivalent for arbitrrary imperfect information games $G$:

(a) $G$ satisfies the commutation properties for Perfect Recall
(b) all instances of axioms (1), (2) are true at every node in G, no matter how we interpret the free propositional variables.

Here, the ‘free’ variables are the ones that do not stand for special game structure such as ‘turn’ or ‘win’. We leave the verification of this standard modal correspondence to the reader. Incidentally, the main purpose of this example is the correspondence method as such, not the particular outcome. E.g., for notions of Perfect Recall without total knowledge of the preceding turn-taking pattern, a similar analysis is possible.

Digression Two ways of thinking

One can think of the above diagrams as constraints passing only bona fide game trees for Perfect Recall. Alternatively, one can allow any imperfect information game, but then use the diagrams as a *PR algorithm* for pruning the tree of dotted lines, leaving only those uncertainties that players with perfect memory will in fact encounter:

*(1)* *remove all uncertainty links that do not originate*
*in an immediately dominating node.*
*(2)* *if no role switch occurred, also remove all uncertainty*
*links between outcomes of different actions*

Analyzing this algorithm, or just the above diagrams, we find another interesting feature of Perfect Recall. Each uncertainty link x…y for a player **i** has a uniquely identifiable *cause* higher up. Working upward along the paths from the root to x, y , one finds some last node where they diverged (at which the other player moves) and a pair of actions taken there which lead to an uncertainty link for **i** . For instance:

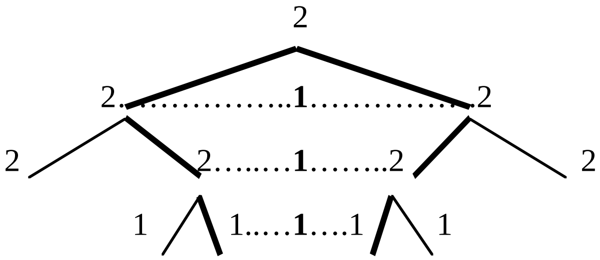

Fig. 10

Thus, *our own uncertainty is always caused by actions of other players that we cannot distinguish*. This general point will return in Section 4.4. The same analysis applies to other phenomena. E.g., axioms so far say nothing about 'forward behaviour'. Thus, Perfect Recall does not exclude that ignorance has a 'miracle cure', by just playing on:

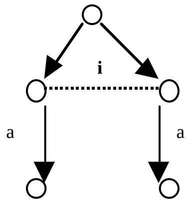

Fig. 11

This is reasonable for context-dependent actions a like "finding out where you are" – but it seems strange with playing 'exactly the same move'. Of course, one might assume that at end points, players *learn* that the game is over. But this would really be one more game move, perhaps in the form of a signal by Nature as an extra player, that should be represented explicitly. In any case, if we require that *playing identical moves propagates the existing uncertainties*

we again make an interchange principle valid, but this time the other way around:

$$[a]\, K_i\phi \rightarrow K_i[a]\phi$$

Knowledge/action interchange principles of both kinds have been studied in temporal action logic (Halpern & Vardi 1989) under the names of "No Forgetting" and "No Learning". In particular, such extras can have a bad effect on the *complexity of validity* in these logics. The technical reason is this. When commuting diagrams enforce a 'grid-like' structure on our models, modal logics can potentially encode so-called Tiling Problems, and thereby even become undecidable. More informally, then, the logic of 'well-behaved' players may be more complex than that of arbitrary players.

**4.3 Information update** The preceding was all about abstract action types $a$ . One can also look at specific cognitive game action such as questions or announcements, and see what knowledge or uncertainty they produce, either publicly, or just for subgroups of players. There is much research on information update in groups of agents, starting from Fagin et al. 1995, and continuing with, amongst many others, van Benthem 1996, Gerbrandy 1998, Baltag 1999, van Ditmarsch 2000. In particular, in addition to keeping track of factual information, *upgrading* of uncertainty relations may affect what we know about each other, and it may affect different players differently:

> "Suppose Nature has dealt me a red card (stating that proposition $\phi$ holds) or a blue one (stating that $\neg\phi$ holds) – say, it was in fact the red card. After this first move, none of us can distinguish the two situations $\phi, \neg\phi$ . Now I look at my card, without showing you the result: but you see what I am doing. Thus I learn what my card is, you still don't know, but you have learnt that I know now. Next, I tell you. Then $\phi$ becomes common knowledge among us."

Let's suppose that I had the red card. Here is a game tree, from which one can read off our successive epistemic states under the moves described, indicated in black:

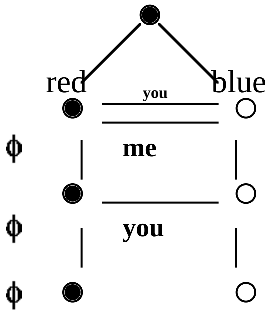

Fig. 12

At the end, both of us have also learnt the actual run of the game. Generalizing from such examples, Baltag 1999 proposes a general update mechanism for cognitive actions in groups that applies very well to games. Its first idea is this. Instead of just having epistemic indistinguishability between states, we can also have plausible

> *epistemic indistinguishabiltiy between actions*.

This made sense with Perfect Recall, where all dotted lines were eventually traceable to other players' actions. This version of epistemic logic is more general than just using dotted lines between outcome states, as different actions can have the same outcomes. So, we include uncertainty between actions as a primitive in our game models, say, as an *equivalence relation* (Baltag update also works more generally). There is much to be said about encoding game actions in this way, such as 'making announcements' (to all or to some players), 'reading signals' (private or public), etcetera. But here, we just give a formal illustration of the basic update mechanism.

> <u>Product Update</u> *Take a game with current state x , and uncertainty relations* $\sim_i$ *among the nodes at x's tree level computed so far. Let a move be made. The new states are the nodes at the next level of the game tree, which can be identified with ordered pairs (*previous state, action last made*). Now we set*
>
> $(y, b) \sim_i (z, c)$ iff $y \sim_i z$ *and* $b \sim_i c$

Thus, new uncertainty equals 'old uncertainty + indistinguishable actions'. This rule produces the intuitively correct results in many cases. E.g., the above envelope example works in this way, with two actions "see red" and "see blue" (which can only be performed succesfully when their content is true), which I can distinguish, but you cannot. Here is an illustration of the effects of product update on general game trees:

<u>Example</u> <u>Updates during play: propagating ignorance along a game tree</u>

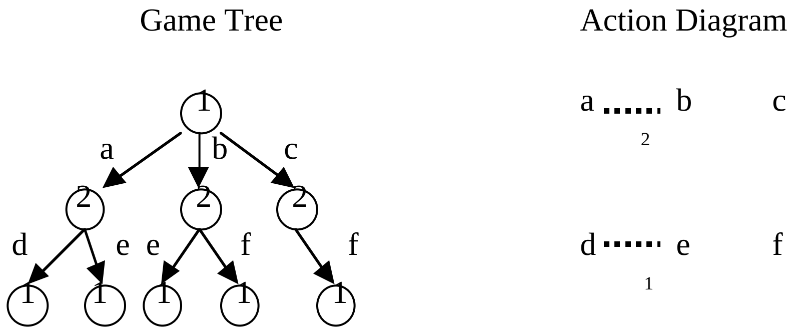

Fig. 13A

Here are the successive updates:

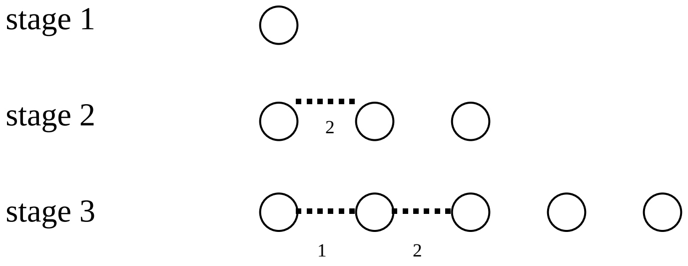

Fig. 13B

**4.4 Update mechanisms and logical axioms** Again, we can analyze such update mechanisms by means of their general dynamic-epistemic properties. For product update, we find three major ingredients with a dynamic-epistemic 'reflection':

(**a**) First, product update implies *Perfect Recall*: the earlier commutation diagrams both follow from the 'product rule', if a player can tell his own actions apart.

(**b**) Next, product also *Propagates Uncertainty* (there are no 'miracles'):

*if* $x \sim_i y$ *, then after performance of* $a$ *in both cases,* $(x, a) \sim_i (y, a)$

With further **i**–indistinguishable actions, we can also allow different a, b here.

(**c**) But there is also a third feature, that one may call *Uniformity:*

Actions are either always distinguishable, or never:

if $(x, a) \sim_i (y, b)$ , then, whenever $u \sim_i v$, also $(u, a) \sim_i (v, b)$

provided the latter moves can be performed at all.

Proposition The three mentioned properties determine product update.

Proof In one direction, we have already shown that the imperfect information trees produced by product update satisfy the three mentioned properties. Conversely, suppose that we have a game tree with all three conditions satisfied. Then we can retrieve an action diagram as follows. Set $a \sim_i b$ if it ever happens in the tree that two nodes $x \sim_i y$ are followed by a, b resp. ending in new nodes that are still $\sim_i$ –related. Then, starting from the root, the $\sim_i$ –links in the tree are precisely those that would be constructed by product update. The reason is this. Perfect Recall says that every uncertainty link in the given tree comes from an uncertainty link one level down, while Propagation says that uncertainty links are inherited downward by indistinguishable actions. Taken together, this is precisely the effect of product update. **QED**

These properties again correspond to special dynamic-epistemic axioms that must hold at every stage of the relevant games. (**a**) The first are the axioms for Perfect Recall. For (**b**), we had the converse axiom $[a]\, K_i\phi \rightarrow K_i[a]\phi$ . Finally, for (**c**), we need a slight extension of the modal language with *universal modalities* $E\phi$, $U\phi$ stating that $\phi$ holds at *some* world, at *all* worlds, resp. (cf. Blackburn, de Rijke & Venema 2000):

$$E(<a^{\smile}>T \;\&\; \neg K_i\neg <b^{\smile}>T) \rightarrow U(<a^{\smile}> \neg K_i\neg\phi \;\rightarrow\; K_i[b^{\smile}]\phi)$$

Thus, also product update corresponds to a special dynamic-epistemic game logic.

**4.5 Playing with bounded memory** At the other extreme of the cognitive spectrum, so to speak, lie ways of playing games with bounded memory (cf. Osborne & Rubinstein, Ch. 9). Suppose that players *forget* all earlier moves, up to $k$ last moves. Perfect Recall then disappears, but the above style of dynamic-epistemic analysis of the resulting game structures still applies. Consider the simplest case of 1-memory, where only the last move is remembered. The dynamics now becomes as follows. At every new stage, for both players, we join all nodes that resulted from the same last action. Then in modal logic terms – again using an auxiliary universal modality, the following axiom will typically be valid on the resulting structures:

$$E(<a>T \;\&\; \phi) \rightarrow U[a^{\smile}]\neg K_i\neg\phi$$

It will be clear how to extend this modal analysis to the case of a $k$-cell memory. This gives a modal take on the considerable body of research in game theory and computer science on strategies that use finite-state machines as their memory.

**4.6 Statics and dynamics again** The above is just a demonstration of the dynamic-epistemic language at work, not a systematic study of game phenomena. But we think it can address such broader issues as the watershed between players' deficiencies and 'objective' ignorance by definition of the game, or the dynamics of further information passing mechanisms such as 'signalling'. But in doing this, the full system would eventually have to combine a *dynamic update logic* for cognitive actions with our more traditional dynamic-epistemic descriptions of the resulting states.

**4.7 Appendix: the influence of modeling decisions** In order to apply dynamic-epistemic logic, one needs a 'modeling phase' for the real phenomena, mapping them to formal actions and uncertainty links. This slack must be kept in mind when judging the framework: formal models can 'fit' to reality in different ways. For instance, actions in real games can be described at different levels of generality. In a card game, a player may “show a card”, which is a generic action, instantiated by concrete actions of showing a specific card. Such decisions may again influence which actions we consider indistinguishable for a player. When showing you one of my cards without looking at it myself, I am uncertain between various concrete show-actions (you are not) – but I am not uncertain about my showing you a card as opposed to doing something else. Thus, the above diagram for Perfect Recall will hold for the latter action, though not for the former. Nice examples of modeling decisions occur in ‘knowledge games’ (van Ditmarsch 2000) with cards, questions, and show-actions.

## 5 Power equivalence

Having dealt with the action level of imperfect information games, we will also briefly look at the more global outcome level, looking for analogues of Sections 2.3, 2.4.

**5.1 Uniform strategies, powers and representation** First, the earlier definition of players’ powers may be adapted in an obvious way:

> *at each node, a player can force those sets of outcomes that*
> *are produced by following one of her* uniform *strategies.*

Example Diminished powers in imperfect information games
Consider Figures 1 and 2. In the former, player **A** has 2 strategies, and player **E** has 4, producing the following set powers:

| | |
|---|---|
| **A** | {1, 2}, {3, 4} |
| **E** | {1, 3}, {1, 4}, {2, 3}, {2, 4} |

With imperfect information added for **E** , **A**’s powers are not affected, but **E**’s are: her 2 *uniform* strategies only give her two set powers:

| | |
|---|---|
| **A** | {1, 2}, {3, 4} |
| **E** | {1, 3}, {2, 4} |

This may be seen as a weakness – but it also shows the much greater modelling power of behaviour by imperfect information games. We can see this by extending the earlier representation result of Section 2.3. First note that the conditions of *superset closure* (C1) and *consistency* (C2) are still valid for powers in imperfect information games.

This may be seen in the above list. What fails there is *determinacy*, as we saw above: **A** does not have the power {1, 4} , but neither does **E** have its complement {2, 3}.

Proposition Any two finite families of sets satisfying conditions (C1), (C2) can be realized as the powers in a two-step imperfect information game.

Proof Instead of stating the procedure formally, we do one illustrative example displaying all the necessary tricks. Suppose that we are given:

| | |
|---|---|
| minimal powers for player **A** | {1, 2, 3}, {3, 4} |
| minimal powers for player **E** | {2, 3}, {1, 4} |

We start with player **A** and provide a number of successor nodes for the power sets, this time with possible duplications, at which player **E** gets to move. The relevant actions here may also involve various duplications. First, we take action types for each set of player **E** , making sure that these get represented via uniform strategies. There may still be 'excess' outcomes in the powers for **E** which we need to 'dilute' by copying and permuting so that they end up in supersets of {2, 3}, {1, 4}:

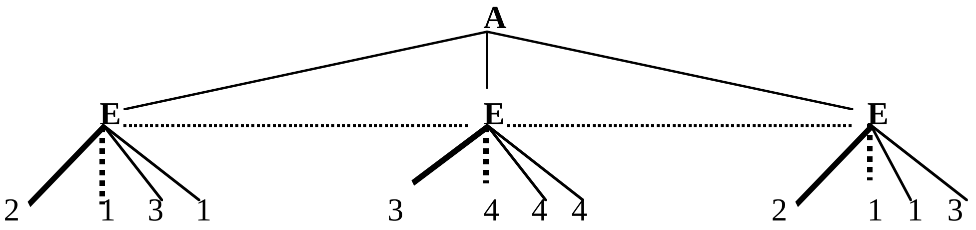

Fig. 14

The sets for the third and fourth uniform strategy are supersets of {2, 3}, {1, 4} .

A further complication arises when **A**'s sets involve outcomes not mentioned in the minimal list for **E** . Say we have:

| | |
|---|---|
| minimal powers for player **A** | {1, 2, 3}, {3, 4} |
| and those for player **E** | {2, 3}, {1, 4, 5} |

Then we need to 'dilute' still more, using similar tricks – this time, involving additional **A**-moves to cases where **E** gets to chose between 1, 4, 5 in such a way that the resulting set powers of **E** are all extensions of {1, 2, 3}, {3, 4} .

Here is the more precise procedure. (**a**) First make sure that all atomic outcomes in players's sets occur for each of them, by adding 'redundant' supersets if needed. Then (**b**) create a preliminary branching for **A** making sure all of its sets are represented. (**c**) Now make sure each outcome set for **E** gets represented as a uniform strategy, by choosing enough branchings from the middle level, starting from the left. This may involve duplication of nodes, as illustrated by the simple case **A** {1, 2}, **E** {1, 2} . (**d**) In case there are outcomes 'left over', repeat the following routine. Suppose outcome i occurring at some mid-level point x was not needed in step (**c**) so far. Then fix any outcome set for **E** produced by some uniform strategy $\sigma$ . Say, the latter chose outcome j at x . Then duplicate this node x , and now add *two new* branchings throughout (for all mid-level nodes considered so far) to obtain two further uniform strategies: one choosing j at x , and i in its duplicate, the other doing the opposite – both following $\sigma$ at all other nodes. This makes outcome j appear as it should for **A** , generating only a harmless superset of outcomes for **E** . **QED**

As in the perfect information case, this result suggests a *logical calculus of game equivalences* for finding these normal forms, which should be a generalization of ordinary Boolean propositional logic. We defer this matter until Section 5.2.

**5.2 Game transformations** The above provides a logic angle to the 'Thompson transformations' of Osborne & Rubinstein, Ch. 11, that transform 'normal form equivalent' imperfect information games into each other. We take them one by one.

Example 1 *Addition of a Superfluous Move* = Idempotence

Figure 206.1 in O&R typically tells us that we can tolerate a new uncertainty by duplicating a move, switching between game trees like the following:

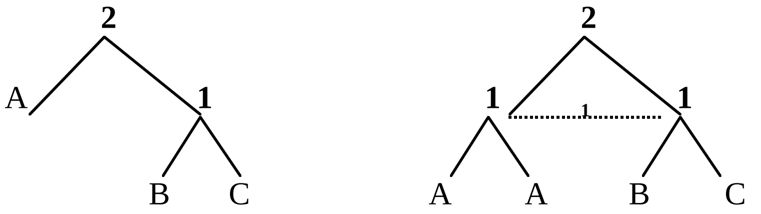

Fig. 15

Computing power relations, we see that nothing changes from left to right at the root. The duplication on the right does not add new outcomes for player **2** , while the two available uniform strategies for player **1** still capture the same two outcome sets {A, B}, {A, C} . This picture is precisely like the propositional Idempotence law

$$A \wedge (B \vee C) \leftrightarrow (A \boldsymbol{\bigvee} A) \wedge (B \boldsymbol{\bigvee} C)$$

where the larger bold face disjunctions indicate the new uncertainty link.

Of course, we are into a non-standard use of propositional logic now, where duplications may matter, and operators can be ‘linked’. In a sense, finite game trees with imperfect information *are* a natural extension of propositional formulas!

Example 2 *Interchange of Moves* = A Deviant Form of Distribution

Figure 208.1 in O&R tells us how to interchange moves of players. The essential transformation here is really one between our earlier Figures 2 and 4 . In both cases, the set powers for the players are (compare the calculation in Section 5.1):

**A** {1, 2}. {3, 4}
**E** {1, 3}, {2, 4}

This is a sort of non-standard propositional ‘Distribution law’:

$$(A \boldsymbol{\bigvee} B) \wedge (C \boldsymbol{\bigvee} D) \quad \leftrightarrow \quad (A \boldsymbol{\bigwedge} C) \vee (B \boldsymbol{\bigwedge} D)$$

with the bold face indicating relevant dotted links. For readers familiar with Hintikka-Sandu IF-games, this is also an equivalence between extended first-order formulas

$$\forall x \ \exists y/\mathbf{x} \ Rxy \ \leftrightarrow \ \exists y \ \forall x/\mathbf{y} \ Rxy$$

Next, we consider the remaining two Thompson transformations:

(a) Inflation-Deflation adds a link that will not materialize in play.
This reflects Perfect Recall, and is one of the moves stated in Section 4.2.

(b) Coalescing Moves says that the zone where a player is to move can be re-encoded as just one bunch of choices. This reflects the emphasis on powers only – disregarding the particular moves, and their order, that players make when it is their turn. This merely states a choice for an ‘outcome-oriented’ notion of game equivalence (rather than a more ‘action-oriented’ one, cf. Section A.1), and has nothing to do with imperfect information as such.

Summarizing our analysis, we have this

Proposition Powers of players in imperfect information games remain the same under the Thompson transformations.

As stated before, we take this to mean that the Thompson transformations are really a kind of complete calculus for a generalized propositional logic. From a more logical point of view, the same effect might be achieved by some modification of ‘informationfriendly' games e.g., in the style of the 3-valued approach of Hintikka & Sandu 1997. Also of logical interest is the *occurrence* character of the above rules. One occurrence of A need not have the same effect any more as several occurrences, when dealing with uniform strategies. This is reminiscent of linear logic, whose game semantics typically inolves non-determined games (Abramsky & Jagadeesan 1994).

Question *Reformulate the complete calculus of finite imperfect information games as a form of propositional logic.*

Appendix To emphasize the analogy once more, we look at the sequence of transformations in O&R, pp. 210/211. It starts with a perfect information game:

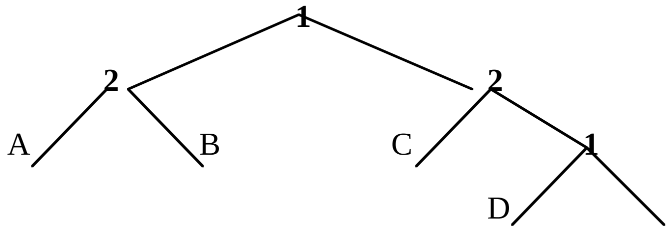

Fig. 16

Here, by the computation of Section 1.4:

powers for player **1** are {A, B}, {C, D}, {C, E}
and for player **2** {A, C}, {B, C}, {A, D, E}, {B, D, E}

By standard propositional transformations as in Section 2.3, this has a distributive normal form which is still a perfect information game:

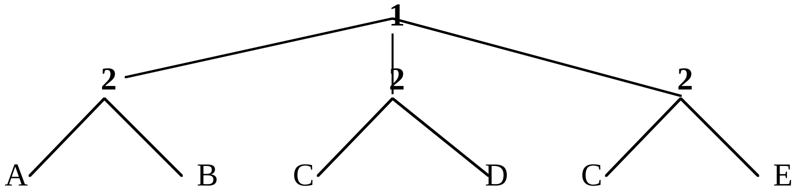

Fig. 17

All steps displayed in O&R preserve powers, too, according to the definition given in Section 2.3, but they result in another normal form:

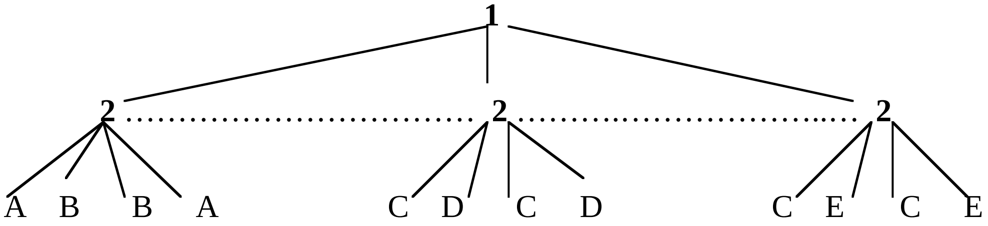

Fig. 18

This analysis is very close to Thompson's original result that his transformations are precisely those that leave the reduced normal form of a game unchanged.

**5.3 Language, uniform strategies and knowledge** Our main approach in this paper has high-lighted the language of epistemic dynamic logic. The latter's 'structural' counterpart for comparing games is epistemic-dynamic bisimulation. But in this section we also considered a coarser game equivalence: up to players' powers for controlling outcomes. Closing the circle, is there also a natural linguistic counterpart to power equivalence? The answer is positive. We need a straightforward epistemic version of the forcing language of Section 2.4 with its 'strategy modalities'

> $\{G, \mathbf{i}\}\phi$
> *player* $\boldsymbol{i}$ *has a uniform strategy forcing a set of outcomes satisfying* $\phi$

The present epistemic setting hides a subtlety, however, concerning the outcome sets. For instance, suppose a player is in an end node $x$ satisfying $p$, but cannot distinguish this from another end node $y$ satisfying $\neg p$. Does the uniform strategy 'do nothing' force $p$ in the end node $x$? Yes, if we look at the actual outcomes produced (just $x$), but no, if we think of the outcomes that seem relevant to the player's deliberations, being both $x$ and $y$. The same point may be made in terms of *knowledge*. If the truth of $\{G, \mathbf{i}\}\phi$ at game state $s$ is to imply knowledge for player $\mathbf{i}$ at $s$ of the eventual forced proposition $\phi$, then we need the latter broader notion of 'relevant outcomes'. We will use the latter notion of uniform $\mathbf{i}$-strategy here. This amounts to computing outcome sets allowing some *third player* arbitrary jumps across uncertainty links for $\mathbf{i}$.

We conclude by comparing this notion to our action language. Can we define this strategy modality in dynamic-epistemic logic, which should be the richer formalism? For *perfect information* games, the answer is positive. The 'fixed-point equation' of Section 2.1 for winning positions can easily e extended to define those game positions from where a player has a '$\phi$–strategy'. With imperfect information games, however, the situation is more delicate. Dynamic-epistemic logic can express subtle distinctions. Do players need to know at each stage that the current strategy is taking them to a desired outcome state (in the above extended sensen of relevant outcomes)? Let us call strategies with this special epistemic property *predictive* (cf. Section 1.2 for its original motivation). To proceed, let us at least assume that we have *Perfect Recall*.

Proposition In imperfect information games with Perfect Recall, all uniform strategies are predictive.

Proof Recall that uncertainties arise in the first place because of moves by the *other* player that you could not distinguish somewhere higher up in the game tree. Then show that the uniform strategy modality satisfies the following fixed-point equation:

$$<G, \mathbf{i}>\phi \quad \leftrightarrow$$
$$((\text{end} \ \& \ K_E \ \phi) \vee (\text{turn}_E \ \& \ \bigvee_a K_E[a]<G, \mathbf{i}>\phi) \vee (\text{turn}_A \ \& \ \mathbf{\&}_a \ K_E[a] <G, \mathbf{i}>\phi)$$

Here the subscript 'a' ranges only over available actions at the relevant game state. In checking this, the direction from left to right is easy. Conversely, one will need to show, in particular, that uniform strategies emanating from different information sets at the current level can be 'patched' to give one uniform strategy at the current node.

One way of testing this equivalence is as follows. As stated above, the relevant moves are now 'game actions plus uncertainty jumps'. Thus, we could also use the original forcing analysis for *perfect information* games, treating combinations of those jumps and the original moves as extra game moves. Note that then, the above epistemic-dynamic combinations $K_E[a]$ are precisely the needed modalities $[\sim_E ; a]$ in that case. In this setting, Perfect Recall says it does not matter in which order we perform the new combined move. Alternatively, we could do the fixed-point analysis by including the moves of the above Third Player explicitly. **QED**

We have not been able to settle the converse question:

*Are all predictive strategies uniform?*

Even if they are broader, or sometimes incomparable vis–à–vis uniform strategies, predictive strategies seem an interesting class of behaviours in their own right.

## 6 Conclusion

We have taken a look at imperfect information games as models for a dynamic-epistemic logic describing actions and outcomes. A systematic analysis like this may help in finding coherent scenarios for the various imperfect information phenomena described in Osborne & Rubinstein 1994. It can also be applied to the much-discussed logical 'IF-games' of Hintikka & Sandu, as shown in van Benthem 2000A. Finally, it suggests a number of more general ways in which logic and game theory can be related.

Our next most pressing business is extending this account is to real games with players' preferences indicated, aspiring to the level of formal sophistication in Stalnaker 1991. In another direction, however, our approach may be seen as an extension of process logics in computer science to deal with imperfect information, something which is needed to really understand modern 'computing systems' such as Internet activities.

## 7 References

S. Abramsky & R. Jagadeesan, 1994, 'Games and Full Completeness for Multiplicative Linear Logic', *Journal of Symbolic Logic* 59:2, 543 – 574.

A. Baltag, 1999, 'A Logic for Communication'. Centre for Mathematics and Computer Science, Amsterdam.

J. Barwise & L. Moss, 1997, *Vicious Circles*, CSLI Publications, Stanford.

P. Battigalli & G. Bonanno, 1999, 'Synchronic Information, Knowledge and Common Knowledge in Extensive Games', *Research in Economics* 53:1, 77–99.

P. Blackburn, M. de Rijke & Y. Venema, 2000, *Modal Logic*, Cambridge University Press, Cambridge.

G. Bonanno, 1992, 'Players' Information in Extensive Games', *Mathematical Social Sciences* 24:1, 35–48.

G. Bonanno, 1993, 'The Logical Representation of Extensive Games', *International Journal of Game Theory* 22, 153–169.

B. de Bruin, 2000, *Modeling Knowledge in Games. Topology and Logic*, Master's thesis, ILLC, University of Amsterdam.

R. Fagin, J. Halpern, Y Moses & M. Vardi, 1995, *Reasoning about Knowledge*, The MIT Press, Cambridge (Mass.).

J. Gerbrandy, 1998, *Bisimulations on Planet Kripke*, dissertation, ILLC, University of Amsterdam.

J. Halpern & M. Vardi, 1989, 'The Complexity of Reasoning about Knowledge and Time', *Journal of Computer and Systems Science* 38:1,195-237.

J. Hintikka & G. Sandu, 1997, 'Game-Theoretical Semantics', in J. van Benthem & A. ter Meulen, eds., *Handbook of Logic and Language*, Elsevier, Amsterdam, 361–410.

R. Moore, 1985, 'A Formal Theory of Knowledge and Action', Research Report, SRI International, Menlo Park.

M. Osborne and A. Rubinstein, 1994, *A Course in Game Theory*, MIT Press, Cambridge.

R. Parikh, 1985, ‘The Logic of Games and its Applications’,
*Annals of Discrete Mathematics* 24, 111–140.
M. Pauly, 2001, *Logic and Games*,
dissertation, ILLC, University of Amsterdam.
R. Stalnaker, 1999, ‘Extensive and Strategic Form: Games and Models for Games’,
*Research in Economics*, 53:293-291.
J. van Benthem, 1996, *Exploring Logical Dynamics*,
CSLI Publications, Stanford.
J. van Benthem, 1998, ‘Dynamic Odds and Ends’,
ILLC Tech Report ML-1998–08, Amsterdam.
J. van Benthem, 1999, *Logic and Games*, electronic lecture notes,
Amsterdam & Stanford, http://www.turing.uva.nl/~johan/
J. van Benthem, 2000A, ‘Hintikka Self–Applied’,
to appear in L. Hahn, ed., *Library of Living Philosophers*.
J. van Benthem, 2000B, ‘When are Two Games the Same?’,
Talk at LOFT-III, Torino 1998, ILLC Tech report X-2000-03, Amsterdam.
H. van Ditmarsch, 2000, *Knowledge Games*,
dissertation, ILLC, University of Amsterdam & Groningen.