

## OPEN PROBLEMS IN LOGICAL DYNAMICS

Johan van Benthem, Amsterdam & Stanford, December 2005

### *Abstract*

In recent years, a number of 'dynamic epistemic logics' have been developed for dealing with information, communication, and interaction. This paper is a survey of conceptual issues and open mathematical problems emanating from this development.

### **1 Logical Dynamics**

The traditional paradigm of logic is drawing a conclusion from some given premises. But derivation from data already at our disposal is just one way in which information can be obtained. We can also observe new facts, or just ask some better-informed person whom we trust. Concomitantly with all this information flow, our knowledge and beliefs change, and this adaptation process may even be triggered by further cues. Such cognitive actions are of logical interest per se, and their explicit study and its various repercussions has been described as a 'Dynamic Turn' in logic (van Benthem 1996). In particular, relevant actions in this broader setting need not be single-agent tasks such as drawing a conclusion or observing some fact. After all, perhaps the simplest logical scenario for getting or giving information is *asking a question*. But this essentially involves information flow between two agents, and their mutual epistemic and 'social' interactions as the question is asked and an answer is given.

An excellent framework for multi-agent dynamic behaviour in communication is *epistemic logic* (introduced in Section 2), suitably 'dynamified' by using ideas from the *dynamic logic* of actions. Section 3 is about the best-explored system of this kind, viz. the dynamic logic of *public announcements*. Section 4 generalizes this to general *dynamic-epistemic logic* of actions or events whose observation conveys information. The resulting technical questions here blend into issues about more classical logical systems, which are discussed in Section 5 on first-order and fixed-point logics. But knowledge and ignorance are not the only attitudes of participants in a conversation. They also have beliefs about their current situation and expectations about the future. These are revised as observation and communication take place. Thus, epistemic dynamics runs into belief revision. Section 6 is devoted to links between dynamic-epistemic logic and *belief revision theory* as developed in AI and related areas.

Next, there is also a longer term to information flow beyond individual update or revision steps. For, these do not occur in isolation. There is a history of past interactions, determining our trust in our interlocutor, as well as a future of things yet to be said. Eventually, this calls for a merge of epistemic and a *temporal logic* which allows for statements of regularities over time, often in the form of 'protocols' (Section 7). Another longer-term perspective concerns the *purpose* of a question. Behind every ordinary question there is a "Why" question: what is the point, and what are people trying to achieve? This leads to current *game logics* for describing strategic behaviour in games, of which there exist quite a few by now. These are just some of the current bridges between logic and game theory, which deserve a separate treatment. Section 8 provides a very brief introduction with pointers. Finally, Section 9 summarizes a few general issues playing across all of the above phenomena.

A word of clarification. 'Open problems' in new areas like this may be of several different sorts. Some are clear-cut mathematical questions, where all notions have crystallized out, and what remains to be done is some hard-nosed *theorem proving*. But other significant questions concern *mathematical modeling*, finding perspicuous formal mechanisms for information update, or temporal evolution, or say, the logical behaviour of different types of agent. Often, questions like this are triggered by the challenge of describing some communicative practice, or some given game. Other questions at this conceptual level have to do with relating different logical paradigms trying to describe the same phenomena. Third, and finally, there are interesting and highly non-trivial questions of *computational implementation* and *cognitive realism* in studying the fit between all these logics and actual behaviour of men and machines. In this survey, the main emphasis will be on the first type of question, but there are a few of the second kind as well. We only mention issues of the third kind in passing.

Finally, a caveat. This paper is not a self-contained introduction to logics for epistemic update and games. It is rather intended for readers with at least some background in the area, who can then see a coherent panorama of directions to be pursued.

## 2 Standard Epistemic Logic

The basis for all our further topics is standard epistemic logic, created originally by Hintikka for analyzing philosophical issues in epistemology – but linked more closely with computer science, and even economics, for a long time now. Excellent introductions to epistemic logic with a computational slant are Fagin, Halpern, Moses & Vardi 1995, van der Hoek & Meijer 1995. For modal logic in general, which serves as a sort of mathematical laboratory, see Blackburn, de Rijke & Venema 2001.

## 2.1 Language

*Definition 1* Standard language.

The standard syntax of epistemic logic has a propositional base with modal operators  $K_i\phi$  (' $i$  knows that  $\phi$ '),  $C_G\phi$  (' $\phi$  is common knowledge in group  $G$ '):

$$p \mid \neg\phi \mid \phi \vee \psi \mid K_i\phi \mid C_G\phi$$

We write  $\langle i \rangle \phi$  for the dual modal existential statement  $\neg K_i \neg \phi$ : which says that agent  $i$  considers  $\phi$  possible. The dual of  $C_G\phi$  is written  $\langle C_G \rangle \phi$ . ♣

*Example 1* Questions and answers.

Let  $Q$  ask a factual question " $P?$ ", where  $A$  answers truly: " $Yes$ ". A presupposition for giving a normal truthful answer is that  $A$  knows that  $P$ :  $K_A P$ . The question itself, if it is a normal co-operative one, also conveys its own presuppositions, such as (i)  $\neg K_Q P \wedge \neg K_Q \neg P$  (' $Q$  does not know if  $P$ ') and (ii)  $\langle Q \rangle (K_A P \vee K_A \neg P)$  (' $Q$  thinks it possible that  $A$  knows the answer'). After the whole two-step communication episode,  $P$  has become common knowledge among  $Q, A$ :  $C_{\langle Q, A \rangle} P$ . Note the crucial role of epistemic iterations: knowledge that agents have about each others knowledge or ignorance, and also the 'group knowledge' achieved at the end. ♣

Another group notion is 'distributed knowledge'  $D_G\phi$ , which holds intuitively when agents in  $G$  put their information together. More generally, epistemic logics can be extended by strengthening their operators in many ways, just as in modal logic in general (cf. Blackburn, van Benthem & Wolter, eds., to appear).


## 2.2 Semantics

*Definition 2* Models and truth definition.

Models  $M$  for the language are triples  $(W, \{\sim_i \mid i \in G\}, V)$ , where  $W$  is a set of worlds, the  $\sim_i$  are binary accessibility relations between worlds, and  $V$  is a propositional valuation. The major epistemic truth conditions are as follows:

$$\begin{aligned} M, s \models K_i\phi & \text{ iff } && \text{for all } t \text{ with } s \sim_i t: M, t \models \phi \\ M, s \models C_G\phi & \text{ iff } && \text{for all } t \text{ that are reachable from } s \text{ by some} \\ & && \text{finite sequence of } \sim_i \text{ steps } (i \in G): M, t \models \phi \end{aligned}$$

*Example 2*



In the black world, the following are true:

$$\begin{aligned}
 &P, \quad K_A P, \quad \neg K_Q P \wedge \neg K_Q \neg P, \quad K_Q(K_A P \vee K_A \neg P) \\
 &C_{(Q, A)}(\neg K_Q P \wedge \neg K_Q \neg P), \quad C_{(Q, A)}(K_A P \vee K_A \neg P) \quad \clubsuit
 \end{aligned}$$

Common knowledge acts as the dynamic logic modality  $[(\bigcup_{i \in G} \sim_i)^*]\phi$ . Finally, distributed knowledge involves intersection of accessibility relations:

$$\mathbf{M}, s \models D_G \phi \quad \text{iff} \quad \text{for all } t \text{ with } s \bigcap_{i \in G} \sim_i t: \mathbf{M}, t \models \phi \quad \clubsuit$$

### 2.3 Basic model theory

A basic notion states when two epistemic models represent the same informational situation from the viewpoint of our basic language.

*Definition 3* Epistemic bisimulation.

A *bisimulation* between epistemic models  $\mathbf{M}, \mathbf{N}$  is a binary relation  $\equiv$  between states  $m, n$  in  $\mathbf{M}, \mathbf{N}$  such that, whenever  $m \equiv n$ , then (a)  $m, n$  satisfy the same proposition letters, (b1) if  $m R m'$ , then there exists a world  $n'$  with  $n R n'$  and  $m' \equiv n'$ , (b2) the same 'zigzag clause' holds in the opposite direction. ♣

Every model  $(\mathbf{M}, s)$  has a smallest bisimilar  $(\mathbf{N}, s)$ , its 'bisimulation contraction'. The latter is the simplest representation of the epistemic information in  $(\mathbf{M}, s)$ .

*Fact 1* Invariance for bisimulation.

For every bisimulation  $E$  between two models  $\mathbf{M}, \mathbf{N}$  with  $s E t$ ,  $s, t$  satisfy the same formulas in the epistemic language with common knowledge.

*Theorem 1* Epistemic definability of models.

Each finite model  $(\mathbf{M}, s)$  has an epistemic formula  $\delta(\mathbf{M}, s)$  (with common knowledge) such that the following are equivalent for all models  $\mathbf{N}, t$ ,

- (a)  $\mathbf{N}, t \models \delta(\mathbf{M}, s)$ ,
- (b)  $\mathbf{N}, t$  has a bisimulation  $\equiv$  with  $\mathbf{M}, s$  such that  $s \equiv t$ .

For a proof, cf. van Benthem 1997, Barwise & Moss 1997. Thus there is a strongest epistemic assertion one can make about states in a current model. In particular, each world in a bisimulation contraction has a unique epistemic definition inside that model. ♣

Also, it follows that, at least when comparing finite models, epistemic equivalence of worlds amounts to the existence of a bisimulation connecting them. For infinite models, the situation is more complex, but this technical line of research is not relevant to us here. Bisimulation is the basic structural equivalence between epistemic models. It plays the same role as potential isomorphism in first-order logic, and like the latter, it has Ehrenfeucht-Fraïssé-style game versions. Many meta-properties of first-order logic, such as interpolation or preservation theorems, transfer to basic epistemic logic without common knowledge by bisimulation-based arguments. For the complicating role of common knowledge, cf. Section 4 below.

*Remark 1* The language with distributed knowledge is richer. Standard modal bisimulations do not preserve statements with intersection modalities, such as  $\langle R \cap S \rangle p$ .



#### 2.4 Axiomatics

The complete logics for these models are the usual ones (cf. the cited references). Minimal modal **K** arises for knowledge if nothing is required of the accessibility relations. In that case, the *K*-operator is better read as *belief*. But most often in this paper, we think of accessibilities as equivalence relations, in which case the complete logic is multi-**S5**. The axioms for  $C_G\phi$  resemble those of dynamic logic, and can be found in the literature. The axioms for  $D_G\phi$  also resemble standard modal ones.

#### 2.5 Complexity

Model checking of epistemic formulas in a given finite models takes polynomial time **P**, just as for modal formulas in general. The complexity of the satisfiability problem for epistemic formulas is **NP-complete** in the case of a single agent, but with two or more agents, it becomes **Pspace-complete**, just like for the minimal modal logic – and even **Exptime-complete** when common knowledge is added to the language.

#### 2.6 Open problems, even here

Static epistemic logic may seem a closed chapter of research. But there are clear open ends, related to alternatives for the above semantics. Here are three of many examples. (a) The universal quantifier truth condition for knowledge has an existential quantifier competitor, taking knowledge of  $\phi$  as the *existence of compelling evidence* for  $\phi$ . Thus, epistemic logic might include an explicit calculus of evidence and reasons (cf. van Benthem 1993). One such system is the 'logic of proofs' of Artemov 1994,

though it has no account of groups and epistemic actions yet. (b) A second alternative are *topological models*, where a universal modality  $K\phi$  describes the interior of the set of points in a model verifying  $\phi$ . Van Benthem & Sarenac 2004 show that this gives much more freedom of epistemic expression. E.g., unlike relational semantics, topological models separate the different intuitions concerning common knowledge distinguished in Barwise 1988. (c) Finally, knowledge in natural usage also takes more object-like arguments, such as "know who did it", "know his telephone number" or "know how to do the job". Such settings are even much more central to epistemic logic in philosophy. No compact useful systems of this sort have been developed, though there is, of course, 'epistemic predicate logic' doing part of the job.

But now, let us move where we really want to go: the 'dynamic turn' putting actions that convey and change information at centre stage.

### 3 Public announcement: epistemic logic dynamified


The basic paradigm of an information-changing action is saying something in public. In the simplest case, someone says something she knows, and this occurs in public, well-understood by all members of the relevant group. The resulting epistemic dynamics has been studied since around 1990. It is a surprisingly rich pilot case for the true realities of everyday communication, where people say things on weaker grounds, with a certain amount of hiding and secrecy. In this section, we give basic definitions and known results, and then discuss a number of interesting questions that have remained open till to-day. For further background in dynamic logic of programs, cf. Harel, Kozen & Tiuryn 2000. The specific open questions discussed here continue an earlier broader survey in van Benthem 2002B.

#### 3.1 World elimination : the system PAL

An answer "Yes" to a question " $P?$ " is a public announcement of the proposition  $P$  in the group  $\{Q, A\}$ . Such an announcement changes the initial model.

*Example 3* Answering a question.

With the above model, announcing  $P$  would just leave the black world:



To the right,  $C_{(Q,A)}P$  holds.



Thus, public announcement involves *world elimination*:

*Definition 4* For any model  $M$ , world  $s$ , and formula  $P$  true at  $s$ ,  $(M|P, s)$  (' $M$  relativized to  $P$  at  $s$ ') is the submodel of  $M$  whose domain is  $\{t \in M \mid M, t \models P\}$ . ♣

This simple process describes many epistemic puzzles and scenarios in the literature. We have a universe of information states (static epistemic models) related by possible transitions: viz. announcement actions  $P!$  taking models  $M$  to submodels  $M|A$ .

*Definition 5* *PAL* language and semantics.

The language of *public announcement logic PAL* is the above epistemic language, with added action expressions

Formulas	$P:$	$p \mid \neg\phi \mid \phi \vee \psi \mid K_i\phi \mid C_G\phi \mid [A]\phi$
Action expressions	$A:$	$P!$

The semantic clause for the dynamic action modality is as follows:

$$M, s \models [P!]\phi \quad \text{iff} \quad \text{if } M, s \models P, \text{ then } M|P, s \models \phi \quad \clubsuit$$

Actually, there are two natural kinds of model here. One is the universe of *all* epistemic models. But smaller natural models can also be 'conversation spaces' with just some initial model plus all its updates by successive 'admissible assertions'.

Many notions from standard epistemic model theory apply in this dynamic setting. E.g., public announcement *respects bisimulation* in the sense of van Benthem 1996:

If  $M, s$  has a bisimulation with  $N, t$ , then, for any epistemic assertion  $P$ ,  
 $M|P, s$  has a bisimulation with  $N|P, t$ .

Restricting the given bisimulation to these two submodels is a bisimulation between the updated models. As a result, the language of *PAL* is invariant for epistemic bisimulations in the earlier sense.

Statements like  $[P!]K_i\phi$  or  $[P!]C_G\phi$  state what agents know after announcements. Thus, the complete logic for *PAL* is an exact mathematical calculus of communication.

*Theorem 2* *PAL* without common knowledge is axiomatized completely by the usual axioms for epistemic logic plus the following *reduction axioms*:

$$[P!]q \quad \leftrightarrow \quad P \rightarrow q \quad \text{for atomic facts } q$$

$$\begin{aligned}
[P!]\neg\phi &\leftrightarrow P \rightarrow \neg[P!]\phi \\
[P!]\phi \wedge \psi &\leftrightarrow [P!]\phi \wedge [P!]\psi \\
[P!]K_i\phi &\leftrightarrow K_i[P!]\phi
\end{aligned}$$

These axioms provide an obvious reduction procedure taking any dynamic *PAL* formula to an equivalent one in static epistemic logic. It follows that

*Corollary 1* *PAL* is decidable.

As it stands this translation is exponential, but Lutz 2004 proposes a more efficient one for *SAT* purposes, claiming that satisfiability for *PAL* remains *Pspace*-complete. Incidentally, model checking for *PAL* is in *P* – as observed by several people.

The full language of *PAL* with common knowledge raises some complications. There is no obvious reduction axiom for assertions  $[P!]C_G\phi$ . This issue has been resolved only in Kooi & van Benthem 2004. First, an extension is needed of the basic epistemic language, with a new modality of *relativized common knowledge*

$$\mathbf{M}, s \models C_G(\phi, \psi) \quad \text{iff} \quad \psi \text{ holds after every finite sequence of accessibility steps for agents going through } \phi\text{-worlds only.}$$

Then we do have a valid equivalence

$$[P!]C_G\phi \leftrightarrow C_G(P, [P!]\phi)$$

Relativized common knowledge is not definable in the basic epistemic language – but it is bisimulation-invariant, and existing completeness proofs are easily generalized. On this extended base, we have the following valid general reduction axiom in *PAL*:

*Theorem 3* *PAL* with relativized common knowledge is axiomatized completely by adding the reduction law  $[P!]C_G(\phi, \psi) \leftrightarrow C_G(P \wedge [P!]\phi, [P!]\psi)$ .

Finally, the dynamic character of *PAL* can be taken a step further. Conversation usually involves sequences of assertions, governed by regular *program constructions* as in dynamic logic. These include

- (a) sequential composition ;
- (b) guarded choice *IF ... THEN... ELSE...*
- (c) guarded iterations *WHILE... DO...*


*Example 4* The puzzle of the Muddy Children.

Here is a simple story that occurs in many variants:


*After playing outside, two of three children have mud on their foreheads. They all see the others, but not themselves, so they do not know their own status. Now their Father comes and says: “At least one of you is dirty”. He then asks: “Does anyone know if he is dirty?” The children answer truthfully. As this question–answer episode repeats, what will happen?*

Nobody knows in the first round. But then, the muddy children both know their status, as each of them can argue as follows. “If I were clean, the one dirty child I see would have seen only clean children around her, and so she would have known that she was dirty at once. But she did not. So I must be dirty, too!” This is symmetric for both muddy children – so both know in the second round. The third child knows it is clean one round later, after they have announced that. The puzzle easily extends to more clean and dirty children. Here is the update sequence for this particular case:

Updates start with the Father's public announcement that at least one child is dirty. This is about the simplest communicative action, and it merely eliminates those worlds from the initial model where the stated proposition is false. I.e., *CCC* disappears:



When no one knows his status, the bottom worlds disappear:



The final update is to *DDC \**

Clearly, the conversation instructions involve all three program operations. ♣

It is easy to add such operations to create compound assertions for more complex conversational instructions. *PAL* must then be extended with the usual axioms for *PDL*. But the combination is surprising, as has been shown by Miller & Moss 2003:

*Theorem 4* *PAL* with *PDL* operations is undecidable.

Indeed, Muddy Children involves a form of parallel program composition as well, since children speak simultaneously in each round. No canonical *PAL* system has been proposed yet for dealing with this further complexity of communication.

This ends our lightning tour of *PAL*. It might seem we know all there is to this system, but this is far from true! Here are four excursions into the unknown.

### 3.2 What are the real update laws?

The earlier complete axiom system leaves many questions unanswered.

**Generalized models** Consider restricted families of epistemic models with assertions  $P$  only admissible as long as the new model  $M|P$  is still inside the family – the above 'conversation models'. Let some assertion  $P$  be forbidden in such a model. Then the principle  $P \rightarrow \langle P! \rangle T$  fails, though it follows from the above *PAL* axioms.

What is the complete *PAL* logic of all generalized models of this sort?

There maybe a simple relativization trick at work here. E.g., consider the reduction axiom for knowledge, in its valid existential form on standard models:

$$\langle P! \rangle \langle i \rangle \phi \leftrightarrow P \wedge \langle i \rangle \langle P! \rangle \phi$$

E.g., from left to right, this is invalid on our general models, since the announcement action  $P!$  available in the current situation  $(M, s)$  need not be available at some shifted situation  $(M, t)$  with  $s \sim_i t$ . But we can put in safeguards, reformulating the axiom to a valid version which assumes that the relevant actions are present:

$$\langle P! \rangle \langle i \rangle \phi \wedge \langle i \rangle \langle P! \rangle T \leftrightarrow \langle P! \rangle T \wedge \langle i \rangle \langle P! \rangle \phi$$

We put  $\langle P! \rangle T$  for the availability of the announcement action  $P!$ , not just its earlier precondition  $P$ . We still have  $\langle P! \rangle T \rightarrow P$  – but as we saw, not its converse.

**Structural rules** Here is a more dynamic notion of inference in models with update (Veltman 1997, van Benthem 1996). Conclusion  $\phi$  follows *dynamically* from  $P_1, \dots, P_k$  if, after public announcements of the successive premises, all worlds in the new information state satisfy  $\phi$ . In terms of *PAL*, the following implication must be valid:

$$[P_1!][\dots][P_k!] C_G \phi$$

This notion behaves differently from standard logic in its structural rules. Permutation of premises, contraction of the same premise, or monotonicity adding premises all fail.

*Theorem 5* The structural properties of dynamic inference are axiomatized by:

*Left Monotonicity*  $X \Rightarrow A$  implies  $B, X \Rightarrow A$

*Cautious Monotonicity*  $X \Rightarrow A$  and  $X, Y \Rightarrow B$  imply  $X, A, Y \Rightarrow B$

*Left Cut*  $X \Rightarrow A$  and  $X, A, Y \Rightarrow B$  imply  $X, Y \Rightarrow B$

This completeness extends to a modal language of all structural properties of announcements over the universe of all epistemic models (cf. van Benthem 2003C). But it remains to be determined how this substructural logic of announcement actions relates to substructural calculi motivated by resource management (Restall 2000).

***Schematic validities*** Unlike most standard logical calculi, *PAL* update logic is not *substitution-closed*. Most conspicuously, its basic axiom

$$[P!]q \leftrightarrow (P \rightarrow q)$$

for atoms  $q$  fails when we substitute arbitrary formulas  $\phi$  for  $q$ . After all, the point of update is that truth values of complex assertions like "I don't know if  $\psi$ " can change.

*Definition 6* Substitution core.

The *substitution core* of update logic consists of those schemata in the language of *PAL* all of whose substitution instances are valid formulas.



There are interesting principles valid in this sense, which did not surface in the earlier 'formula by formula' axiom system for *PAL*. An example is iterated announcement:

*Fact 2*  $[A![[B!] \phi \leftrightarrow [[A!]B!] \phi$  is a valid schematic principle of *PAL*.

Here is a total list of schematically valid principles for public announcement which covers all cases we have been able to find so far. Again, we take existential versions:

$$\begin{aligned} & [P!]T \\ [P!]\perp & \leftrightarrow \neg P \\ [P!]\neg\phi & \leftrightarrow P \rightarrow \neg[P!]\phi \\ [P!](\phi \wedge \psi) & \leftrightarrow [P!]\phi \wedge [P!]\psi \\ [P!]K_i\phi & \leftrightarrow K_i[P!]\phi \end{aligned}$$

$$\begin{aligned}
[P!]CG(\phi, \psi) &\leftrightarrow CG(P \wedge [P!]\phi, [P!]\psi) \\
[A!][B!]\phi &\leftrightarrow [([A!]B)!]\phi
\end{aligned}$$

*Question 1* Is the substitution core of *PAL* decidable; or at least axiomatizable?

### 3.3 Model theory of learning

Understanding the effects of communicative actions can be highly non-trivial. E.g., it might seem 'obvious' that public announcement of a proposition  $P$  in a group results in common knowledge of  $P$ . But there is no axiom  $[P!]C_G P$  in the complete list for *PAL*. And indeed, the intuition fails. Here is a much-cited counter-example.

*Example 5* Self-refuting assertions.

Let  $p$  be the case, but I don't know this. If you know both these facts, and announce truly "You don't know that  $p$ , but it is true" ( $\phi = \neg K_{you} p \wedge p$ ), the current model gets updated to one with  $p$  true throughout, and  $p$  becomes common knowledge between us. But then the announced statement  $\phi$  becomes false by its very announcement.



Other examples are the ignorance assertions in the puzzle of the 'Muddy Children' whose repeated public announcement eventually led to common knowledge in the group. This observation raises an issue of model-theoretic preservation.

*Question 2* Which syntactic shapes of formulas  $P$  are 'self-fulfilling', i.e.: they guarantee common knowledge of  $P$  after their announcement?

Van Benthem 2002B notes that self-fulfilment holds for the epistemic equivalent of 'universal formulas', true in any submodel of models where they hold:

*Fact 3* All shapes generated by the following grammar are self-fulfilling:

$$(\neg)p, (\neg)q, \dots \mid \phi \wedge \psi \mid \phi \vee \psi \mid K_i \phi \mid C_G \phi$$

Here the final clause allows arbitrary common knowledge formulas. But there are other self-fulfilling cases, at least when the accessibility relations in our models are equivalence relations. An example guaranteeing common knowledge is then  $\neg K p$ .

*Remark 2* Most general postconditions.

Van Benthem 2002B relates the difficulty of the above preservation question to the

impossibility of defining 'most general postconditions', in the sense of computer science, for assertions inside *PAL*. Stating a fact  $p!$  has a most general postcondition  $C_G p$ , but the argument is ad-hoc. In general, the generic description of the strongest effect of  $P!$  is backward-looking: "there *was* a model from which the current one arose by announcing  $P$ ". This requires *temporal* past operators beyond *PAL*. ♣

Self-fulfilling propositions arise in many ways.

*Example 6* Verificationism and learnability.

Van Benthem 2004C discusses the Verificationist Thesis in epistemology, which holds that "all true assertions may be known". This calls for assertions which make some given truth common knowledge. There are several relevant versions of self-fulfilment then, ranging from more local to more global. Validity of  $[\phi!]C_G\phi$  means that a true statement  $\phi$  can always be learnt by announcing  $\phi$  itself. But a statement  $\phi$  is learnable in a weaker sense if there is some assertion  $\psi$  such that  $\phi \rightarrow [\psi!]C_G\phi$  is valid. And it is learnable in a still weaker sense if that trigger  $\psi$  may depend on the model where  $\phi$  is true. These notions form a strict hierarchy in models for multi-*S5*, where they are all are all decidable. It is an open question whether this decidability of learnability notions extends to other model classes for epistemic logic, say *S4*. ♣

Thinking more locally in a given model, van Benthem 2002B does show that

*Fact 4* In finite models, any announcement with a proposition has an update which can be generated equivalently by a proposition which becomes common knowledge after its announcement.

No version is known for this local persistence result on infinite models.

### 3.4 *Communication and planning*

*PAL* is not just a formalism for analyzing given statements or longer conversations. It can also be used for planning assertions, just as dynamic logics of programs can be used for both analysis and design. Here is one obvious issue that arises then.

**Maximal communication** Here is what a group can achieve by maximal public announcement. Epistemic agents in a model  $(M, s)$  can tell each other things they know, thereby cutting down the model to smaller sizes, until nothing changes.

*Theorem 6* Each model  $(M, s)$  has a unique minimal submodel reachable

by maximal communication of known propositions among all agents.

Up to bisimulation, its domain is the set  $COM(\mathbf{M}, s) = \{t \mid s \cap_{i \in G} \sim_i t\}$ .

An agent  $j$  can even reach  $COM(\mathbf{M}, s)$  by speaking just once, if she takes bisimulation contractions of updates all along the way. For then, the strongest proposition known to  $j$  corresponds to the set  $\{t \mid s \sim_j t\}$ . As this is definable, she can state that definition.  $COM(\mathbf{M}, s)$  is related to the earlier notion of distributed knowledge  $D_G\phi$ . Still, there remains a difference between evaluating the assertion  $\phi$  in the whole model  $\mathbf{M}$ , or just within the submodel  $COM(\mathbf{M}, s)$ . The latter notion seems a better candidate for the intuitive notion of 'implicit group knowledge', but it has not yet been axiomatized.

The importance of relation intersection here suggests another open problem.

*Problem 1* Find a natural extended notion of bisimulation

that is characteristic for epistemic logic enriched with  $D_G\phi$ .

**Planning and reachability** More delicate planning issues for assertions include announcing facts publicly between some agents while leaving some others in the dark.

*Example 7* 'Moscow Puzzle'.

7 cards are distributed among  $A, B, C$ .  $A$  gets 3,  $B$  gets 3,  $C$  gets 1. How should  $A, B$  communicate publicly, in hearing of  $C$ , to find out the real distribution of the cards while  $C$  does not? Solutions depend on numbers of cards (van Ditmarsch 2002). ♣

The logic of hiding in public view has some initial observations so far, but no general theory. Of course, 'public privacy' goes only so far, and we will look at real hiding phenomena in communication in Section 3. Generalizing from this example, one can look at any model  $(\mathbf{M}, s)$  (perhaps inside some 'conversation model'), and ask which models  $(\mathbf{N}, t)$  satisfying certain desiderata are *reachable* from  $(\mathbf{M}, s)$  by means of finite sequences of available assertions. Here is one technical notion related to this.

*Definition 7* Learnability.

The *learnability modality* is defined by a quantification over possible assertions:

$\mathbf{M}, s \models \langle \text{learn} \rangle \phi$  iff there is some assertion  $P$  with  $\mathbf{M}, s \models \langle P! \rangle \phi$ . ♣

*Question 3* Is *PAL* plus the operator  $\langle \text{learn} \rangle$  still decidable?

### 3.5 Group knowledge

Epistemic logic is mostly about individual agents and their interaction. But plural agents can also perform actions, like winning a soccer match or electing a president –

and likewise, it makes sense to ascribe knowledge to plural subjects, such as groups or organizations. Common knowledge was one step in this direction, although its definition is essentially still a reduction to knowledge for individual group members. Less reductive was the notion of implicit group knowledge, which arises only if individuals cooperate and share information. One obvious way of doing this is by imagining plural epistemic subjects with internal communication channels, so that subgroups can exchange information (cf. van Benthem 2002B):

*Problem 2*     Develop a version of dynamic-epistemic logic of groups  
with communication channels as primitive entities.

In natural language, we often switch between 'individual' and 'collective' agents – with some predicates applying only to the former, and others only to the latter. Some group predicates are straightforward lifts of individual behaviour for all members ("the boys had the flu"), whereas others are not ("the prisoners liberated each other"). Thus, there are no total reductions between the two levels, and both seem essential in our way of describing the world. The same duality might apply to epistemic agents, and hence, an epistemic logic of collective agents and group-forming operations seems of interest.

#### **4     Dynamic Epistemic Logic**

Many forms of communication have private aspects, inadvertently, or with deliberate hiding. People read cards in public view without showing them to others, they whisper in crowded lecture theatres, and they send each other secret messages along a channel they believe to be safe. Civilized social life is full of procedures where information flows in restricted ways. Such forms of communication are high-lighted in parlour games, which have been designed to manipulate information flow. Cf. van Ditmarsch 2000 for a complete analysis of informational moves in the popular game *Cluedo*. Such moves can be explained to a large extent in *dynamic-epistemic logic*, which provides a principled account of update of information models by 'event models'. The main source for this calculus is Baltag, Moss & Solecki 1998, and a forthcoming textbook is van Ditmarsch, van der Hoek & Kooi, to appear.

##### ***4.1 Information from arbitrary events: product update***

The information models provided by epistemic logic have a natural companion.

*Definition 8*     Event models.

The set of relevant events  $A$  in a communication scenario forms an *event model*

$$\mathbf{A} = (A, \{\sim_i \mid i \in G\}, \{PRE_a \mid a \in A\}),$$

with agents' uncertainty relations  $\sim_i$ . The latter encode which events agents cannot distinguish. E.g., when I read my card, and you know it must be either red or black, you cannot distinguish the two events of 'my reading red' and 'my reading black'. But you can distinguish either from 'my reading orange' or 'Mount Etna erupting'. Finally, actions  $a$  always have *preconditions*  $PRE_a$  for their being executable. We will assume that preconditions are formulated within our language. E.g., 'my reading red' presupposes that I hold a red card – or, as in an earlier example, my asking a question may presuppose that I do not know the answer, but think that you do. ♣

Here is our general way of computing a new information state.

*Definition 9* Product update.

Let models  $(\mathbf{M}, s)$  and  $(\mathbf{A}, a)$  be given. The *product model*  $(\mathbf{M} \times \mathbf{A}, (s, a))$  has domain

$$\{(s, a) \mid s \text{ a world in } \mathbf{M}, a \text{ an action in } \mathbf{A}, (\mathbf{M}, s) \models PRE_a\},$$

and the new uncertainties satisfy

$$(s, a) \sim_i (t, b) \text{ iff both } s \sim_i t \text{ and } a \sim_i b.$$

Thus, new uncertainty can only come from existing uncertainty via indistinguishable events. Finally, the valuation for atoms  $p$  at  $(s, a)$  is copied from that at  $s$  in  $\mathbf{M}$ . The new actual world is the pair  $(s, a)$  of the old actual world and the actual event. ♣

This mechanism was proposed in Baltag, Moss & Solecki 1998, building on earlier work by Gerbrandy, Groeneveld, and Plaza. It can model a wide range of phenomena, including games (van Ditmarsch 2000, van Benthem 2001). In particular, it can deal with misleading actions as well as truthful ones – though this requires leaving the realm of epistemic models with equivalence relations. As with public announcement, truth values of propositions can change drastically under product update.

*Definition 10* Update language and semantics.

The dynamic-epistemic language for this new setting is:

$$p \mid \neg\phi \mid \phi \vee \psi \mid K_i\phi \mid C_G\phi \mid [A, a]\phi$$

where  $(\mathbf{A}, a)$  is any finite event model with actual event  $a$ . Semantic interpretation takes place as in the above, with the new clause

$$\mathbf{M}, s \models [A, a]\phi \quad \text{iff} \quad \mathbf{M} \times \mathbf{A}, (s, a) \models \phi \quad \clubsuit$$

*Theorem 7* Dynamic epistemic logic is effectively axiomatizable and decidable.

Here is the key reduction axiom extending the one for public announcement:

$$[A, a] K_i \phi \leftrightarrow PRE_a \rightarrow \& \{K_i (PRE_b \rightarrow [A, b]\phi) \mid b \sim_i a \text{ in } A\}$$

As before, a reduction axiom for common knowledge in subgroups requires a language extension. Van Benthem, van Eijck & Kooi 2004 do this in a somewhat baroque (but natural) extension of epistemic logic, which allows any *PDL*-style program with tests and regular operations inside the epistemic language. There are also quite different reformulations of dynamic-epistemic logic. E.g., Baltag, Coecke, & Sadrzadeh 2004 provide a very general *co-algebraic* version, relying heavily on the bisimulation invariance of the above language. But in this Section, we concentrate on other issues, having to do more with new features of the product update framework.

#### 4.2 Update evolution

Unlike updates for public announcement, product update can blow up the size of the initial input model  $M$ . Here is a simple illustration.

*Example 9* Blow-up of information models.

Suppose a public announcement of the true fact  $P$  takes place in a group  $\{1, 2\}$ , but 2 is not sure whether it was an announcement of  $P$ , or just some identity event  $Id$  which could happen anywhere. In that case, a two-world model with worlds for  $P$  and  $\neg P$  turns into a three-world model with states  $(p, "P!")$ ,  $(p, Id)$  and  $(\neg P, Id)$ .



A typical example of blow-up occurs in *games*. Players start from an initial situation  $M$ , say a deal of cards, and the event model  $A$  contains all possible moves that they have – with preconditions restricting when these are available at players' turns. Then the game tree consists of all possible evolutions through the nodes in a tree model:

*Definition 11* Update evolution models.

Let an initial epistemic model  $M$  be given, and an event model  $A$ . Then  $Tree(M, A)$  is the infinite epistemic model consisting of disjoint copies of all successive product update layers  $MxA, (MxA)xA, \dots$




But this infinity can be spurious – as happens in many games, where complexity of information can grow, but then decrease again toward the end game.

*Example 10* Stabilization under bisimulation.


Consider a model with two worlds  $P, \neg P$ , between which agent  $1$  is uncertain, though  $2$  is not. The actual world has  $P$ .



Now a true announcement of  $P$  takes place, and agent  $1$  hears this. But agent  $2$  thinks the announcement might just be a statement "True" which could hold anywhere. The event model for this scenario looks like this:



The next two levels of  $Tree(M, A)$  then become as follows:



But note that there is an epistemic bisimulation between these two levels, connecting the lower three worlds to the left with the single world  $(P, P!)$  in  $M \times A$ . Thus,  $M \times A \times A$  is bisimilar with  $M \times A$ , and the iteration remains finite modulo bisimulation. ♣

Van Benthem 2001 determines the exact conditions under which an extensive game with imperfect information can be represented in this tree format. Moreover, van Benthem 2002B stated a 'Finite Evolution Conjecture' saying that, starting from a given finite  $M$  and  $A$ , the model  $Tree(M, A)$  always remains finite modulo bisimulation. This would imply that some horizontal levels  $M \times A^k$  and  $M \times A^l$  in that tree must be bisimilar, with  $k < l$ .

The Finite Evolution Conjecture was refuted in Sadzik 2004, which views  $Tree(M, A)$  as a dynamical system – and then shows, amongst many other things, that

- (a) The Conjecture holds in single-agent  $S5$ -models.
- (b) The Conjecture fails in some models with two  $S5$ -agents. But it holds for many special cases of such models, e.g., when the epistemic accessibility relations for all agents in the model  $A$  are linearly ordered by inclusion.

Sadzik uses finite pebble games over  $M$  and  $A$  to determine when  $Tree(M, A)$  is finite modulo bisimulation. This is the case iff the 'responding player' in the game has a winning strategy. The computational complexity of this game is still unknown. Also,

the above analysis still leaves open the possibility that large classes of scenarios fall into the Finite Evolution class – e.g., those corresponding to most parlour games.

*Remark 3* Action emulation.

Other new questions about product update focus on the behaviour of action models. Van Eijck, Ruan & Sadzik 2004 introduce a relation of *action emulation* between action models  $A$ ,  $B$  which holds iff  $A$ ,  $B$  produce bisimilar results on all bisimilar inputs  $M$ ,  $N$ . This turns out to be different from a simple notion of bisimulation between action models, and it has an interesting independent characterization. ♣

### 4.3 Questions of language design

It is still somewhat of an open question which language has the best expressive power for dynamic-epistemic logic. Here we just state a number of possible directions which have not yet been fully explored.

Baltag, Moss & Solecki 1998 use a standard epistemic language with a common knowledge modality, but this fails to generate intuitive *reduction axioms* for formulas  $[A, a]C_G\phi$  which are crucial to understanding group communication. As mentioned before, van Benthem, van Eijck & Kooi 2004 do find such axioms in a *PDL*-style extension of epistemic logic. But are there natural languages in between '*DEL*' and their '*E-PDL*' that do the job? Or conversely, going to even richer languages, does the same 'reductive equilibrium' also hold for the full epistemic  $\mu$ -calculus?

Another potential extension has to do with *concurrency* in communication, such as the simultaneous assertion by the Muddy Children of their status. This might involve the same sort of *process-algebraic calculi* for concurrency as those developed in the 1980s in computer science. In fact, product update itself is a typical graph operation of the sort studied in that area, which respects modal bisimulation. What is the connection between dynamic-epistemic logic and process algebra? Cf. van Ditmarsch, van der Hoek & Kooi 2003 for a first approach to parallel update actions.

Finally, the system still has a glaring asymmetry. Information models  $M$  interpret a *language*, but event models  $A$  do not! In particular, preconditions  $PRE_a$  are not formulas that are *true* at  $a$  in  $A$ : they refer to what is to be true in  $M$ . But it is easy to introduce a second epistemic language describing properties of actions or events. It has atomic properties, Boolean combinations, and epistemic modalities such as "for some action indistinguishable from the current one". Van Benthem 1999, ten Cate 2002 show how this simplifies product update in a joint language for information and action models. Van Benthem 2005 relates this to modal languages for more general

product operations, and reductions of truth in a product to truth in its components. But a convincing and usable language of event models per se still has not emerged.

#### **4.4 Extensions of empirical coverage**

These were all questions about dynamic epistemic logic and product update as they stand. But in studying real communication, many further issues will come to the fore that might lead to extensions of the framework. Here are a few examples where no mathematical theory worth reporting has been developed as yet.

- (a) In Section 7, we will look at temporal settings where we can look at the past of some current update process, as well as its future.
- (b) Van Benthem & Liu 2004 show that different update rules hold for *different types of agents*, including those which are not logically omniscient, because of memory limitations. They also suggest a general study of heterogeneous groups, whose members do not all update in the same way. This is like 'bounded agents' in game theory.
- (c) Bleeker & van Eijck 2000 study *security and cryptographic protocols* in a product update model which includes world-changing actions.
- (d) Castelfranchi 2004 points out that belief dynamics goes hand in hand with dynamics of *changing goals and intentions*. The latter dimension has been left out completely in the logicians' systems so far.

Communication is a vast area, with many 'thresholds' of complexity, e.g., when moving from public to private actions, or from plain speaking to misleading or lying. The eventual hope would be that dynamic epistemic logics, of whatever enriched sort, help in locating and understanding these practices.

## **5 Background in standard logics**

Epistemic logic has usually been considered as an applied system, rather than a vehicle for mathematical research as such. And indeed, for technical purposes, it is often easy to look at related systems in which it can be embedded. There are well-known direct translations from epistemic logics into modal or dynamic logics, but also into first-order logic, and into fixed-point logics. We mention some open problems in those areas that seem relevant to our concerns in the preceding sections.

### **5.1 Modal logic**

**Model-changing operators**  $[P!]\phi$  is really an instance of a sort of modal operator that occurs more often in the recent literature. Its evaluation *shifts the current model*. Other examples of such operators are the 'deletion modality'  $\langle - \rangle\phi$  of van Benthem 2002A, which states truth in some submodel with one accessibility link deleted:

$$M, s \models \langle - \rangle\phi \quad \text{iff} \quad \text{for some } (s, t) \in R_M, W_M, R_M - \{(s, t)\}, V_M, s \models \phi$$

This notion seems simple, but Loedding & Rohde 2003 show that this modal logic has a *Pspace*-complete model checking problem, while satisfiability is undecidable. Another example are the 'bisimulation quantifiers'  $\langle bis-p \rangle\phi$  of Hollenberg 1998:

$$M, s \models \langle bis-p \rangle\phi \quad \text{iff} \quad \text{for some model } (N, t) \text{ with a bisimulation} \\ \text{for the language of } \phi \text{ minus } p \text{ between } M, N \text{ linking } s \text{ to } t, N, t \models \phi$$

This evaluates  $\phi$  in some model bisimilar to the current one, disregarding truth values for the atom  $p$ . As an extension of *PDL*, this has equal expressive power with the modal  $\mu$ -calculus, and uses of this formalism are still increasing. The modalities  $\langle A, a \rangle\phi$  of dynamic epistemic logic belong to the same family, and hence their expressive power and complexity behaviour are of interest.

**Generalized consequence** Model-jumping also occurs inside modal logic. In many completeness proofs, one has a modal formula  $\phi$  true in some model  $(M, s)$ , and then finds another model  $(N, t) \models \phi$  with nicer structural properties, bisimilar to  $(M, s)$ . This suggests two new notions of modal consequence:

$$\phi \Rightarrow [bis] \psi \quad \psi \text{ holds in all models bisimilar to a model for } \phi \\ \phi \Rightarrow \langle bis \rangle \psi \text{ every model for } \phi \text{ is bisimilar to one for } \psi$$

Van Benthem 1997 shows that the former notion is recursively axiomatizable and even decidable. For the second, he merely proves its equivalence to conservativity of  $\psi$  over  $\phi$  w.r.t. existential consequences:

$$\text{for all existential modal formulas } \alpha, \text{ if } \psi \models \alpha, \text{ then } \phi \models \alpha$$

*Question 4* Is existential bisimulation consequence decidable, too?

## 5.2 First-order logic

**Relativization** Announcing  $A$  amounts to the well-known first-order operation of *relativizing* a model  $M, s$  to a definable submodel  $M|A, s$ . This may also be described via syntactic relativization of formulas  $\phi$  by the update assertion  $A$ :

$$M|A, s \models \phi \text{ iff } M, s \models (\phi)^A.$$

The reduction axioms of *PAL* are just the usual inductive definition of relativization. The point of our relativized common knowledge was that the epistemic language with just  $C_G\phi$  is not closed under relativization, whereas the basic dynamic logic *PDL* is. Even so, *PAL* is a modal axiomatization of model-theoretic relativization.

**Problem 3** Give a complete logic of relativizations  $(\phi)^A$  in *first-order logic*.

One new valid principle at this level is

$$((A)^B)^C \text{ is logically equivalent to } {}_A((B)^C) \quad \textit{Associativity}$$

This is the first-order counterpart of the earlier valid principle  $[A![[B!]\phi] \leftrightarrow [([A!]B)!] \phi$  in the substitution-closed schematic version of *PAL*.

**Relative interpretability** Relativization often occurs together with other operations on models, such as *translation* of predicates and formation of new objects as *pairs* – e.g., in the notion of relative interpretation of one theory into another. An example is the way in which the theory of the rationals is embedded into that of the natural numbers. Rationals may be viewed as pairs of natural numbers, and one then looks at a definable subset of  $N \times N$ , with some newly defined predicates. Product update involves exactly the same features, and hence dynamic epistemic laws may also be viewed as an axiomatization of such further model-forming operations. Thus, *DEL* might point the way toward a systematic modal meta-theory of model-theoretic operations.

**Interpolation** Update is also related to *interpolation*. Barwise & van Benthem 1999 define a general notion of 'entailment along a relation'  $R$ :  $\phi$  *entails*  $\psi$  *along*  $R$  if,

$$\text{for all models } M, N \text{ with } M R N, \text{ if } M \models \phi, \text{ then } N \models \psi.$$

Reasoning often involves using what we know about one situation to infer properties of *another*. Standard logical consequence is entailment along the identity relation. Of

particular interest in a modal setting is entailment along modal bisimulation. Now we can formulate interesting combined preservation-interpolation theorems.

*Theorem 8* The following are equivalent for all first-order formulas  $\phi, \psi$ .

- (a)  $\phi$  entails  $\psi$  along bisimulation,
- (b) there exists a modal formula  $\alpha$  with  $\phi \models \alpha \models \psi$ .

The modal interpolant  $\alpha$  is the 'bridge' which allows the jump from models  $M$  where  $\phi$  holds to bisimilar models  $N$  where  $\psi$  holds. In the special case where  $\phi$  is invariant for bisimulation (i.e.,  $\phi$  entails itself along bisimulation), this gives the usual result that  $\phi$  must be equivalent to some modal formula. Similar interpolation results might make sense for dynamic-epistemic logic. Here is a simple illustration (recall Fact 4).

*Example 11* Interpolation by self-fulfilling statements?

Consider a valid *PAL* statement  $\phi \rightarrow [P!]\psi$  about the effects of a  $P$ -announcement. Can we always use self-learning statements as 'interpolants' to explain this, in the sense that  $\phi \rightarrow [P!]\psi$  is valid iff there exists some self-fulfilling assertion  $\alpha$  such that (i)  $\phi \models \alpha$ , (ii)  $\alpha \models \psi$ , (iii)  $\models [\alpha!]\mathcal{C}_G\alpha$ ? The answer seems negative. ♣

### 5.3 Fixed-point logics

Fixed-point extensions exist for both modal and first-order languages (cf. Harel, Kozen & Tiuryn 2000, Stirling 1999 on the  $\mu$ -calculus, and Ebbinghaus & Flum 1995 on *LFP(FO)*). These languages contain all predicates which are definable as smallest or greatest fixed-points of monotone set operations. In particular, with epistemic logic, common knowledge may be viewed as a greatest fixed-point

$$\forall p \bullet \phi \wedge \bigwedge_{i \in G} K_i p$$

Many features of epistemic logic with common knowledge become clearer against this background. E.g., the two characteristic principles for complete axiomatization express the two aspects of this definition: an axiom stating the fixed-point equation by itself, plus an induction rule stating that this is a *greatest* fixed-point.

But there is also a less attractive side to this reduction. For, the explicit programs of propositional dynamic logic disappear, and the language becomes 'static' again,

expressing properties of worlds. Nevertheless, programs can still be extracted under certain circumstances. The basic idea of *PDL* is finite *reachability*: each program describes a regular set of traces consisting of basic action steps and test for formulas. Here is a relevant fragment of the  $\mu$ -calculus which contains *PDL*, in which all approximation sequences for smallest or greatest fixed-points stop after  $\omega$  steps.

*Definition 12* The  $\omega$ - $\mu$ -calculus.

The  $\omega$ - $\mu$ -calculus only allows smallest fixed-point operators in the following existential format, then approximation sequences always stabilize by stage  $\omega$ :

$$\mu p \bullet \phi(p) \quad \text{with } \phi \text{ constructed according to the syntax}$$

$$p \mid p\text{-free formulas} \mid \wedge \mid \vee \mid \text{existential modalities.} \quad \clubsuit$$

The reason for guaranteed approximation by stage  $\omega$  is the special syntax of  $\phi(p)$ . Van Benthem 1996 proves a preservation theorem showing the equivalence, for first-order formulas, of this format with the key semantic property of 'Finite Distributivity' for the approximation maps. The  $\omega$ - $\mu$ -calculus is still too strong, though, since even a simple  $\omega$ - $\mu$ -formula like  $\mu p \bullet q \vee (\langle 1 \rangle p \wedge \langle 2 \rangle p)$  is not definable in *PDL*. Still, *PDL* is closed under simultaneous fixed-points of a yet more special syntactic type of recursion, with only disjunctions of existential formulas  $\langle \pi \rangle p$  where the propositional recursion variables  $p$  occurs only in the end position. We omit details (cf. van Benthem, van Eijck & Kooi 2005) – but this provides one more format for reduction axioms in dynamic-epistemic logic. The interesting issue remains which natural fragments of the  $\mu$ -calculus suffice for communication events.

Of the many further questions raised by the connection with fixed-point logics, we point out just one. Evidently, we would like to have analogues of classical meta-properties like interpolation or preservation for our static or dynamic epistemic logics including common knowledge. But it is a bit of a scandal that we do not know if these hold! E.g., interpolation for *PDL* has been open for some 30 years now, with several ship-wrecked proofs in the prestigious published literature. The problem is that the usual model-theoretic arguments based on compactness fail for fixed-point languages. And so do their analogues for infinitary languages (Barwise 1975, Barwise & van

Benthem 1999) – as fixed-point languages strike out from first-order logic in a different way, allowing for an explicit definition of well-foundedness.

*Question 5* Do *PAL* and *DEL* with common knowledge inherit the nice model-theoretic properties of their finitary modal base logics, such as interpolation, Los-Tarski, or Lyndon theorems?

We do not know. The languages seem so simple that a positive answer might be expected, but few proofs of this kind exist. One exception is the propositional  $\mu$ -calculus, where d'Agostino & Hollenberg 2000 established uniform interpolation by *automata-theoretic* methods. These methods might also work here.

## 5 From information update to belief revision

### 5.1 From knowledge to belief

Dynamic-epistemic logic as presented so far seems mainly concerned with knowledge – but this is an artefact of our presentation. One could just as well formulate everything so far in terms of agents' *beliefs*  $B_i\phi$ , interpreted over models for a minimal modal logic without special requirements on the now *directed* accessibility relations. The only technical modification worth pointing out is the issue of 'common belief'. The fixed-point operator for common knowledge enforced veridicality:

$$C_G\phi \leftrightarrow \phi \wedge \&_{i \in G} K_i C_G\phi$$

But the modification is simply this:

$$CB_G\phi \leftrightarrow B_i\phi \wedge \&_{i \in G} B_i C_G\phi$$

Still, there remains a desideratum, since product update does not do genuine *belief revision*. Here is a simple illustration.

*Example 12* Updating with conflicting information.

Suppose that  $p$  is true, but an agent believes that  $\neg p$ :



Now an announcement  $p!$  occurs. The product update rule will eliminate the  $\neg p$ -world, leading to a one-point  $p$ -world with an *empty* accessibility relation. Here, the agent believes a contradiction, or indeed any formula. ♣

Now, there is an easy fix in the preceding case: just make the accessibility relation reflexive in the updated model, and let the agent believe that  $p$ . But the point is that there is no *principled* way of doing this in the current framework. And it seems hard to find a way of modifying accessibility relations dealing with all more delicate cases that arise with as 'strong prior preference for  $\neg p$ ' versus 'strong new evidence for  $p$ '.

## 5.2 *Dynamic doxastic logic*

An extension of product update to deal with belief revision has been proposed in Aucher 2003. The basic idea is as follows. One first introduces a more refined language for belief, with 'graded operators' – an idea due to Spohn 1988:

$B_i^\alpha \phi$  agent  $i$  believes  $\phi$  up to plausibility level  $\alpha$ .

This requires enriched epistemic models  $M$ , retaining the earlier indistinguishability relations  $\sim_i$ , but now expanded with maps  $\kappa_i$  assigning doxastic plausibility values  $\kappa_i(s)$  for agents to worlds  $s$ :

$M, s \models B_i^\alpha \phi$  iff for all  $t \sim_i s$  with  $\kappa_i(t) \leq \alpha$ :  $M, t \models \phi$


This could even be generalized to having world-dependent plausibility values, but there seems to be no need for this in the most basic cases of belief – if we assume that agents know their own beliefs and plausibilities. It is easy to axiomatize the combined epistemic doxastic logic of these models, especially with a trick from Liu 2004. Just add proposition letters  $plaus_i(\alpha)$  to the language, and interpret these as true at just those worlds whose plausibility is at most  $\alpha$ . Then  $B_i^\alpha \phi$  may be defined explicitly as

$K_i(plaus_i(\alpha) \rightarrow \phi)$ ,

and its special properties follow automatically, such as the introspection laws

$B_i^\alpha \phi \rightarrow K_i B_i^\alpha \phi$  and  $B_i^\alpha(B_i^\alpha \phi \rightarrow \phi)$

Next, we lift the same way of thinking to *event models*. These, too, can be viewed as having plausibility structure. E.g., I may hear you say something, and believe that it is most likely to be " $p$ ", although it might also be " $\neg p$ ", as a negation got lost in the wind. Say, in fact, you said " $\neg p$ ": This would give rise to the following event model, where both actions are epistemic possibilities, and one is doxastically preferred:



The preconditions may still encode agents' common knowledge about these actions, which would work exactly as before. One could also make them reflect agents' private beliefs about these actions. E.g., if you think the speaker is a liar, then you would believe that "say  $p$ " tends to happen when  $\neg p$  is in fact the case. The latter sort of modified precondition has not yet been studied – and we continue with the former, simpler case. Quite sophisticated scenarios can occur, even in this idealized format.

In what follows, we omit some complications having to do with the finite range of plausibility values used in Aucher 2003, for technical reasons that are irrelevant here.

*Definition 14* Product update for belief.

Product models  $(M, s) \times (A, a)$  are defined as before for their purely epistemic part. The crucial additional rule updates plausibility values:

$$\kappa_i((s, a)) = \kappa_i(s) + \kappa_i(a) - \text{MIN}(\kappa_i(t) \mid M, t \models \text{PRE}_a\}$$

The correction factor  $\text{MIN}(\kappa_i(t) \mid M, t \models \text{PRE}_a\}$  subtracts the lowest value for worlds in the original model  $M$  satisfying the precondition of the action just performed. ♣

Some drawing of diagrams for simple scenarios will show how this works.

*Example 13* Updating beliefs.

(a) Start from uncertainty about  $p$  while believing that  $\neg p$ , with values 1 for the  $p$ -world, and 0 for the  $\neg p$ -one. Now listen to a true announcement of  $p$ , as in the earlier problematic example of belief collapse after public announcement. Then the plausibility value of the only remaining world ( $p$ , "say  $p$ ") is computed as

$$\begin{aligned} \kappa_i((p, \text{"say } p\text{"})) &= \kappa_i(p) + \kappa_i(\text{"say } p\text{"}) - \text{MIN}(\kappa_i(t) \mid M, t \models p\} \\ &= 1 + 0 - 1 = 0 \end{aligned}$$

This is as it should be: after the update the agent now believes that  $p$ .

(b) Next, let the agent hear a statement which she believes to be a true announcement of " $p$ " (plausibility value 0 in the event model), though it might also be one of " $\neg p$ " (plausibility value 1). Then two worlds remain, and their plausibility values become:

$$\begin{aligned}
\kappa_i((p, \text{"say } p\text{"})) &= \kappa_i(p) + \kappa_i(\text{"say } p\text{"}) - \text{MIN}(\kappa_i(t) \mid \mathbf{M}, t \models p) \\
&= 1 + 0 - 1 = 0 \\
\kappa_i((\neg p, \text{"say } \neg p\text{"})) &= \kappa_i(\neg p) + \kappa_i(\text{"say } \neg p\text{"}) - \text{MIN}(\kappa_i(t) \mid \mathbf{M}, t \models \neg p) \\
&= 0 + 1 - 0 = 1
\end{aligned}$$

♣

Still, these very precise numerical calculations may be more specific than what is intuitively supported by our understanding of belief revision.

*Problem 4* Find a more qualitative version of belief revision after all.

Given this update mechanism for beliefs, it is easy to find complete axiom systems for dynamic doxastic logic in the earlier dynamic-epistemic style. In particular, the product update rule for  $\kappa_i$ -values turns into a reduction axiom for graded beliefs recording the same numerical convention. Still, intuitive questions remain:

The net effect of the product update rule is 'radical': the last observed event largely determines the new beliefs. This is clear in case (b) of the preceding example, where the belief about the statement just heard wipes out prior beliefs, whatever plausibility one would have given before to  $\neg p$ . This is not all we want, however – since our logic must account for the undeniable phenomenon of *different policies for belief revision*: more radical, or more conservative. Liu 2004 studies parametrized variants for different sorts of agent, giving different weights to the factors in the update rule:

$$\lambda \bullet \kappa_i(s) + \mu \bullet \kappa_i(a)$$

*Problem 5* Axiomatize dynamic belief logic with parametrized update rules.

### 5.3 Better-known theories of belief revision

Belief revision theory is mostly associated with the 'AGM' paradigm (cf. Gärdenfors & Rott 1995). The core of this is a system of axiomatic postulates on three operations for changing a current theory  $T$  after some new fact  $A$  comes into focus:

$$\begin{aligned}
&\text{update } T+A, \\
&\text{revision } T*A, \\
&\text{contraction } T-A.
\end{aligned}$$

The third of these is different from what we have looked at so far, giving a theory as much like  $T$  as possible except for leaving out  $A$ . We will continue with the first two.

A semantic account of belief revision was given in Grove 1988. It has close analogies with Lewis-style models for conditional logic, which sets

$M, s \models A \Rightarrow B$  iff  $B$  is true in all  $A$ -worlds that are 'closest' to  $s$ ,  
as measured by some given similarity relation between worlds.

Basically, an update with  $A$  moves a current theory  $T$ , viewed as a set of worlds  $\|T\|$ , to the set of most plausible worlds, as seen from  $T$ 's standpoint, which satisfy  $A$ . This connection between belief revision and conditional logic, suitably dynamified, has been a persistent intuition in the area – with many different more precise formulations. Ryan & Schobbens 1997 is a sophisticated recent example.

*Remark 4* Updates for real change.

Modern belief revision theories include both belief revision concerning one fixed situation, and *world update* in the sense of Katsuno & Mendelzon 1992, where incoming assertions can also report real changes in the world. The latter moves beyond update and revision in the stricter sense of this paper. Nevertheless, it is quite easy to incorporate these notions in the framework of this paper – e.g. by allowing events in action models to change truth values of proposition letters. Cf. van Benthem, van Eijck & Kooi 2004 for one particular implementation. ♣

The setting of *AGM* is not straightforwardly comparable with those for dynamic-doxastic update, since it only considers non-iterated theory change for single agents, and that only on the basis of non-doxastic factual statements. For instance, *AGM* starts with an apparently intuitive 'Success Postulate'  $A \in T^*A$ . But this only makes sense for update with non-doxastic factual statements, witness the earlier discussion of the Learning Problem with public announcement. Moreover, the repertoire of three operations is much smaller than that of the above plausibility event models, which can model infinitely many different action scenarios. And finally, the *AGM* postulates do not provide systematic reduction axioms for beliefs after the update.

But there is an alternative modal-logic based reconstruction of *AGM*-style belief revision which does provide some features closer to this paper. Segerberg 1995 (following the abstract 'dynamic modal logic' of update in van Benthem 1989, de Rijke 1992) proposes a so-called 'dynamic-doxastic logic' *DDL* with *PDL*-style operators

$[+A]\phi$ ,       $[*A]\phi$

for update and revision with non-modal statements. These modal operators satisfy a quite standard modal set of axioms expressing, amongst others, the partial functionality of these operations, and the fact that factual assertions lead to belief. One semantics for this system uses *hyper-theories*, which are like the models for 'premise semantics' of conditionals in Veltman 1985 (going back also to Kratzer and Lewis). Hyper-theories  $H$  are families of sets of worlds, pre-encoding the potential revisions an agent is willing to make. They satisfy some technical mathematical assumptions which we omit here. Here is just an illustration of how they work.

*Example 15* Update and revision of hyper-theories.

- + Update with  $A$  takes a hyper-theory  $H$  to the union of (i)  $H$ , (ii)  $\bigcap H \cap \|\mathbf{A}\|$
- \* Revision with  $A$  takes  $H$  to the union of (i)  $\text{cons}(H, A)$ : i.e., all sets in  $H$  with a non-empty intersection with  $\|\mathbf{A}\|$ , and (ii)  $\bigcap(\text{cons}(H, A)) \cap \|\mathbf{A}\|$  ♣

With these stipulations, *DDL* gets valid reduction axioms in the *DEL* style like

$$[+A] B_i \phi \leftrightarrow B_i (A \rightarrow \phi), \quad [*A] B_i \phi \leftrightarrow (A \Rightarrow_i \phi)$$

where  $A \Rightarrow_i \phi$  is the earlier conditional assertion, interpreted in the Lewis style, but now with 'closeness' as 'relative plausibility' for the relevant agent  $i$ . This is much like the reduction axioms in van Benthem 2003A for belief after public announcement. Girard 2004 explores further analogies between our brand of dynamic doxastic logic and that of the Segerberg systems. But evidently, much remains to be clarified.

#### 5.4 Probabilistic update

Dynamic-doxastic logic also raises another issue, relating to the other broad tradition dealing with information update. Belief update is well-known in the much more established *Bayesian probabilistic* format, where prior probability distributions over propositions are modified to new ones by giving various weights to the priors and the new observation. The first systematic connections with dynamic-epistemic logic have been made in Kooi 2003. A full-fledged product update system for probabilities is found in van Benthem 2003B. Still, this is only the first merge, and this bridge between epistemic logic and probability could be broadened considerably.

*Problem 6* Make a systematic comparison of *DEL*-style and Bayesian update.

Moreover, the connection between plausibility approaches like that of Spohn-Aucher, and probability approaches to belief needs to be understood. E.g., one difference is that plausibility approaches freely combine beliefs, because of the validity of

$$B_i^\alpha \phi \wedge B_i^\alpha \psi \rightarrow B_i^\alpha (\phi \wedge \psi)$$

But this principle is typically invalid in probabilistic update, where probabilities of conjunctions tend to fall below those of their components.

## 6 Temporal Epistemic logic

### 6.1 Broader temporal perspectives on update

Product update in *DEL* modifies knowledge about the present situation. In particular, all uncertainty relations for agents are 'horizontal' inside the current model. But there is a much broader temporal setting, from past to future. For instance, many epistemic puzzles contain dialogues like "I do not know if *P*", "I *knew* that", where the latter past tense refers to the initial state, not the one updated by the first assertion. A technical motivation pointing the same way was the need for temporal past operators when defining strongest *postconditions*, of communicative actions.: cf. Remark 2 above.

Other motivations for looking back at the past of some update process are when we find out that someone has been lying (Yap 2004). This seems to call for some sort of *backward temporal update* – which remains to be defined. Of course, all this requires maintaining a record of previous updates. Such a perspective is natural from many viewpoints. E.g., consider the earlier idea of public announcement as *relativization*. So far, we discarded old information states. But now, we can keep the old state, and perform 'virtual update' via relativized assertions. Thus, the initial state already contains all possible future communicative developments. This is more in line with the above update evolution models  $Tree(\mathbf{M}, \mathbf{A})$  which contain all possible conversational trajectories. Such models obviously support a richer temporal language.

A final motivation for taking a broader temporal perspective comes from more global information that we may have about some communication process. E.g., we might know that, however history unfolds, every question will always be answered. Such 'fairness' or 'liveness' properties can be formulated in standard temporal logics over  $Tree(\mathbf{M}, \mathbf{A})$ , now viewed as a branching-time structure. Such systematic informative restrictions on the possible runs of a system are often called *protocols*.

### 6.2 Knowledge and ignorance over time

Epistemic-temporal frameworks have existed since the 1980s. One famous example is the run-based model for distributed systems of Fagin, Halpern, Moses & Vardi 1995, or the epistemic branching temporal logic framework of Parikh & Ramanujam 2003. Consider models for branching time with nodes  $s$  and histories  $h$  representing runs of a game. As a process – say, a game – proceeds, agents are in a node on some actual history whose past they know, but whose future is yet to be fully revealed. We will think of this setting as a *tree of finite sequences of events*, just as happens in temporal logics in computer science. Sometimes, a selection is made among all possible branches in the tree, leaving just the 'legal runs' obeying the relevant protocol.

One convenient basic language has proposition letters for local properties of nodes, Boolean operations, as well as temporal and modal operators.

*Definition 15* Branching time semantics.

Formulas are interpreted at nodes  $s$  on histories  $h$ , in a format with clauses

$$(a) \quad \mathbf{M}, h, s \models F_a \phi \quad \text{iff} \quad s \frown \langle a \rangle \text{ lies on } h \text{ and } \mathbf{M}, h, s \frown \langle a \rangle \models \phi$$

The standard operator  $F$  ('at some point in the future') is the transitive closure of this one-step modality, taken over all possible events  $a$ .

$$(b) \quad \mathbf{M}, h, s \models P_a \phi \quad \text{iff} \quad s = s' \frown \langle a \rangle, \text{ and } \mathbf{M}, h, s' \models \phi$$

Again,  $P$  ('at some point in the past') is the transitive closure.

$$(c) \quad \mathbf{M}, h, s \models \langle \rangle \phi \quad \text{iff} \quad \mathbf{M}, h', s \models \phi \text{ for some history } h'$$

which coincides with  $h$  up to stage  $s$ .



*Remark 5* Simultaneity.

Language extensions also make sense. In particular, enriching this temporal-modal language helps describe event trees more explicitly. E.g., a sideways modality for simultaneity refers to truth at the end sequences of the same length. ♣

One usually reads the modal operator  $\langle \rangle$  as an absolute historical possibility. But one can also view it as some sort of epistemic possibility for agents, referring to future continuations which they think possible. In the latter case, we should enrich the above models and allow different sets of continuations for different agents.

Zanardo 2002 has an extensive discussions of this temporal setting, possible axioms, etc. In particular, protocols generalize models to a more manageable one allowing for

an axiomatization in a Henkin 'general model' style. Belnap, Perloff & Xu 2001 is a more philosophical-logical analysis of knowledge, belief and conditionals in branching time, focusing on the action logic *STIT* ("see to it that").

*Problem 7* Give a systematic comparison of *STIT* with *DEL/DDL*.


If players  $i$  also have beliefs about the course of the game, then we add binary relations  $\leq_i$  to the model of *relative plausibility*, and we add a doxastic modality

*Definition 15, continued*

(d)  $M, h, s \models B_i \phi$  iff  $M, h', s \models \phi$  for all histories  $h'$  which coincide with  $h$  up to stage  $s$  and are most plausible for  $i$  according to the given relation  $\leq_i$ .

Here is the existential dual version of this:

(d)'  $M, h, s \models \langle B, i \rangle \phi$  iff  $M, h', s \models \phi$  for some history  $h'$  coinciding with  $h$  up to stage  $s$  and most plausible for  $i$  according to the given relation  $\leq_i$ :



This can model different views of players on how the process will develop. ♣

*Problem 8* Find a most suitable extension of the *DEL*-language that works well computationally on these event trees.

### 6.3 Representation of update logics

The above branching-time models specialize naturally to the event-tree semantics that comes with product update. Van Benthem 2001, van Benthem & Liu 2004 observe that product update will produce very special uncertainty patterns in such trees.

*Theorem 10* An event tree is isomorphic to  $Tree(M, A)$  for finite models  $M, A$  iff it satisfies suitable properties of 'Perfect Recall' and 'Uniform No Learning'.

Similar representation theorems are provable for agents which have bounded memory to various degrees. Even so, a general update logic of different sorts of agents, and what happens when they interact, is missing so far. This would be a sort of update counterpart of attempts at giving up 'logical omniscience' in static epistemic logic.

The current 2005 version of van Benthem 2003B reformulates dynamic-epistemic update logics in this broader setting, replacing the earlier reduction axioms for update and revision by more standard temporal-doxastic versions. (Cf. also van der Meiden 1998 on looking for regular subcases of temporal run models where anything goes.)

*Problem 9*     Relate *DEL* to the more general semantics of messages  
                    in branching-time models of Parikh & Ramanujam 2003.

Problems 8,9 are part of a more general task – which also came up in discussions at the 'First Indian Congress on Logic and its Relationship with Other Disciplines', Mumbai 2005. Epistemic-temporal logics are used widely in analyzing and designing computational processes, as well as games in general (cf. van der Meijden 2005, Ramanujam 2005, Venkatesh 2005). One would like to merge the various approaches mentioned here into one framework, and then explore the fine-structure of specific types of agent, message, and general action. *DEL* describes the simplest setting, while in the general framework 'anything goes'. What lies in between?

Interestingly, the latter paper also connects our qualitative sort of message update with quantitative information in the sense of mathematical Information Theory. Van Rooij 2003 makes a similar move in the semantics of natural language.

Once connections like these have been made in some appropriate fashion, it will also be clear how to increase the scope of current update logics to include the earlier-mentioned phenomena of Section 6.1, such as more complex temporal preconditions for actions – instead of just atemporal epistemic ones – , announcements of temporal assertions, and actions that change truth values of atomic propositions. Each such step represents an upward move in expressiveness toward the total system of the two mentioned references. Another interesting topic in this setting are the earlier-mentioned protocols constraining general runs of a communication sequence. Protocol information is absent from *DEL* (or *DDL*) as such. But, through the use of preconditions in event models, *DEL* does rule out certain runs of subsequent events 'locally', so to speak. The question is which types of protocol are needed: these might be classified by syntactic types of definition in the epistemic-temporal language.

#### **6.4 Connections with other parts of mathematics**

The above discussion suggests links between the dynamic logics of information update over time in this paper and a number of established areas from mathematics. These connections are still to be explored in depth. Here are a few examples.

(a) Protocols might be just any set of branches in event trees. But the latter are like Baire spaces with their natural *topology*. Thus, well-behaved sets of branches lie at various levels of the topological Borel Hierarchy over our trees. Still, pure tree topology may not suffice, as the epistemic structure of indistinguishability between tree nodes matters. Then we seem to need generalized notions of topology, as natural closure conditions w.r.t. indistinguishability relations might become relevant.

(b) We already mentioned *probability theory* as a more quantitative account of belief and update, and likewise, *information theory* as a quantitative measure of channel transmission. Unifying such viewpoints seems a major undertaking at this stage.

(c) Event trees are also the natural universe for *learning theory* (Kelly 1996), whose connections with modal and dynamic approaches seem obvious – though just a few initial links have been made so far (Hendricks 2003).

(d) Van Benthem & Liu 2004 suggest connections between the study of different epistemic agents and *automata theory*, e.g., with Nerode representations for trees with added uncertainty relations. This topic is also connected to the use of finite automata for agents with 'bounded rationality' in game theory. A final mathematical connection here is the study of Turing machines with ignorance for players about what is observed on input and output tapes, as in Ann Condon's pioneering work on such devices in computational *complexity theory*.

(e) And finally, our long-term temporal perspective suggests an obvious step from dynamic logic to mathematical *dynamical systems theory*. This move has been made in other settings, too – as one studies long-term bulk behaviour of logical inference.

## 7 Game Logics and Game Theory

As explained in the Introduction to this paper, communication between agents naturally leads to the topic of *strategic interaction* generally. The best current model for such interactions are *games*. Thus, we get into connections between logic and *game theory*. 'Logic and Games' is a fast-developing interface by itself these days, with many different sub-themes, including strong influences from computer science. In particular, there are two strands of research to be distinguished: using ideas from game theory in logic, and using ideas from logic in game theory in logic.

We are currently preparing a separate companion paper (van Benthem 2005A) with a similar list of open questions on Logic and Games, along both of these lines. Here is a

preliminary table of contents, whose headings form a natural continuation of the topics addressed for update and communication in this paper:

- (a) Dynamic logics and strategy calculus in extensive games
- (b) Dynamic logics of players' powers for determining outcomes
- (c) Operations that form new games out of given ones, and game algebra
- (d) Logics of preferences and rational action
- (e) Logics with imperfect information
- (f) Infinite games and temporal evolution in linear and temporal logic
- (g) Fine-structure of game theory: justifying Nash equilibria  
in strategic and extensive games
- (h) Logic games with preferences?
- (i) Probability, expectations and mixed strategies
- (j) Evolutionary games and dynamical systems

Surveys of relevant topics, as well as some concrete results, are in van Benthem 1999–2004, van Benthem 2001, 2002B, 2003C, 2004A. Recent dissertations showing the interest of the interface are Pauly 2001, de Bruin 2004, Dechesne 2005, Harrenstein 2004. The final paper will contain a much more extensive bibliography.

## 8 Conclusion

This paper is a survey of interesting new themes and questions in the dynamic logic of communication. In line with the purpose of the present Volume, most of its open problems concern mathematical *system issues* about dynamic epistemic logics – and *technical junctions* to be made between different areas of logic and mathematics.

Still, it is worth repeating a point from our Introduction about other sorts of issues that are of equal importance. In particular, on the descriptive side, there are valid concerns about the adequacy of the formalizations proposed here. E.g., descriptive frameworks for rational agents in computer science or cognitive science tend to view information update in a broader setting of *beliefs*, *desires* and *intentions* (the 'BDI' paradigm), with a concomitant dynamics of preferences, goals, and plans. Similar broad models have been around in computational linguistics since the 1980s. Sticking with plain update, or even belief revision, may be a beneficial mathematical idealization – but it may also suffer from the 'Inventor's Paradox' of trying too little.

Another missing aspect is *cognitive reality*. Update, communication, and conversation seem typical topics for controlled experiments, but no systematic collaboration between logicians and cognitive scientists has developed yet along these lines.

Finally, staying with logic itself, putting communication and many-agent interaction at centre stage might have far-reaching repercussions for the agenda of the field – and even the things we want to know about its key systems. E.g., which meta-theorems (old, or new) for first-order logic are most relevant to its communicative, rather than its inferential use? We leave these larger issues to the philosophers of logic.

## 9 References

- G. d'Agostino & M. Hollenberg, 2000, 'Logical Questions concerning the  $\mu$ -Calculus: Interpolation, Lyndon & Los-Tarski', *Journal of Symbolic Logic* 65, 310-332.
- S. Artemov, 1994, 'Logic of Proofs', *Annals of Pure and Applied Logic* 67, 29–35.
- G. Aucher, 2003, *A Joint System of Update Logic and Belief Revision*, Master of Logic Thesis, ILLC University of Amsterdam.
- A. Baltag, B. Coecke, & M. Sadrzadeh, 2004, 'Algebra and Sequent Calculus for Epistemic Actions', *ENTCS Proceedings of Logic and Communication in Multi-Agent Systems (LCMAS) Workshop, ESSLLI 2004, Nancy, France*.
- A. Baltag, L. Moss & S. Solecki, 1998, 'The Logic of Public Announcements, Common Knowledge and Private Suspicions', *Proceedings TARK 1998*, 43–56, Morgan Kaufmann Publishers, Los Altos. Many updated versions.
- J. Barwise, 1975, *Admissible Sets and Structures*, Springer, Berlin.
- J. Barwise, 1988, 'Three Views of Common Knowledge', *Proceedings TARK 1988*, 365–397, Morgan Kaufmann Publishers, Los Altos.
- J. Barwise & J. van Benthem, 1999, 'Interpolation, Preservation, and Pebble Games', *Journal of Symbolic Logic* 64:2, 881–903.
- J. Barwise & L. Moss, 1997, *Vicious Circles*, CSLI Publications, Stanford.
- N. Belnap, M. Perloff & M. Xu, 2001, *Facing the Future*, Oxford University Press, Oxford.
- J. van Benthem, 1989, 'Semantic Parallels in Natural Language and Computation', in H-D Ebbinghaus et al., eds., 1989, *Logic Colloquium. Granada 1987*, North-Holland, Amsterdam, 331-375.
- J. van Benthem, 1993, 'Reflections on Epistemic Logic', *Logique et Analyse* 34 (vol. 133-134), 5-14.
- J. van Benthem, 1996, *Exploring Logical Dynamics*, CSLI Publications, Stanford.

- J. van Benthem, 1997, 'Dynamic Bits and Pieces' and Pieces', Report LP-97-01, Institute for Logic, Language and Computation, University of Amsterdam.
- J. van Benthem, 1999, 'Radical Epistemic Dynamic Logic', note for course 'Logic in Games', ILLC, University of Amsterdam.
- J. van Benthem, 1999–2004, *Logic and Games*, electronic lecture notes, <http://staff.science.uva.nl/~johan/>, Amsterdam and Stanford.
- J. van Benthem, 2000, 'Logics for Information Update', *Proceedings TARK VIII*, Morgan Kaufmann, Los Altos, 51–88.
- J. van Benthem, 2001, 'Games in Dynamic Epistemic Logic', *Bulletin of Economic Research* 53:4, 219–248.
- J. van Benthem, 2002A, 'An Essay on Sabotage and Obstruction', in D. Hutter, ed., *Festschrift for Jörg Siekmann*, Springer Verlag.
- J. van Benthem, 2002B, 'Extensive Games as Process Models', *Journal of Logic, Language and Information* 11, 289–313.
- J. van Benthem, 2002C, 'One is a Lonely Number: on the Logic of Communication', Tech Report PP-2002-27, ILLC Amsterdam. To appear in P. Koepke et al., eds., *Colloquium Logicum*, AMS Publications, Providence.
- J. van Benthem, 2003A, 'Belief over Time', Manuscript, ILLC Amsterdam. Current version: Update and Revision in Games', to be presented at APA-ASL Symposium on Games, San Francisco, March 2005.
- J. van Benthem, 2003B, 'Conditional Probability Meets Update Logic', *Journal of Logic, Language and Information* 12:4, 409–421.
- J. van Benthem, 2003C, 'Rational Dynamics and Epistemic Logic in Games', in S. Vannucci, ed., *Logic, Game Theory and Social Choice III*, University of Siena, department of political economy, 19–23.
- J. van Benthem, 2003D, 'Structural Properties of Dynamic Reasoning', in *Meaning: the Dynamic Turn* (J. Peregrin, ed.), Elsevier, Amsterdam, 2003, 15–31.
- J. van Benthem, 2004A, 'A Mini-Guide to Logic in Action', *Philosophical Researches*, Suppl., 21–30, Beijing, Chinese Academy of Sciences.
- J. van Benthem, 2004B, 'What One May Come to Know', *Analysis* 64 (282), 95–105.
- J. van Benthem, 2005A, 'Open Problems in Game Logics', ILLC Amsterdam, to appear, 'Games, Logic, and Computation': <http://www.ilic.uva.nl/lgc/>.
- J. van Benthem, 2005B, 'Two Logical Concepts of Information', to appear in L. Moss, ed., *Memorial Volume for Jon Barwise*, Bloomington, Indiana.
- J. van Benthem, J. van Eijck & B. Kooi, 2004, 'Logics for Communication and

- Change', ILLC Amsterdam, CWI Amsterdam, Philosophical Institute, University of Groningen.
- J. van Benthem & F. Liu, 2004, 'Diversity of Logical Agents in Games', Report PP-2004-13, ILLC Amsterdam. Also in *Philosophia Scientiae* 8:2, 163–178.
- J. van Benthem & D. Sarenac, 2004, 'The Geometry of Knowledge', Tech Report PP-2004-20, ILLC Amsterdam.
- P. Blackburn, J. van Benthem & F. Wolter, eds., to appear, *Handbook of Modal Logic*, Elsevier Science Publishers, Amsterdam.
- P. Blackburn, M. de Rijke & Y. Venema, 2001, *Modal Logic*, Cambridge University Press, Cambridge.
- A. Bleeker & J. van Eijck, 2000, 'The Epistemics of Encryption', CWI Report INS-R0019, Amsterdam.
- B. de Bruin, 2004, *Explaining Games*, Dissertation 2004-03, ILLC Amsterdam.
- Ch. Castelfranchi, 2004, 'Reasons to Believe: Cognitive Models of Belief Change', ISTC-CNR, Roma – Workshop *Changing Minds*, ILLC Amsterdam.
- B. ten Cate, 2002, 'Internalizing Epistemic Actions', in M. Martinez, ed., *Proceedings of the NASSLLI-1 Student Session*, Stanford University, 109–123.
- B. ten Cate, 2005, *Model Theory for Extended Modal Languages*, Dissertation, Institute for Logic, Language and Computation, University of Amsterdam.
- F. Dechesne, 2005, *Game, Set, Maths*, Ph.D. Dissertation, Philosophical Institute, Katholieke Universiteit Brabant, Tilburg.
- H. van Ditmarsch, 2000, *Knowledge Games*, dissertation DS-2000-06, Institute for Logic, Language and Computation, University of Amsterdam & Department of Informatics, University of Groningen.
- H. van Ditmarsch, 2002, 'Keeping Secrets with Public Communication', Department of computer science, University of Otago, Dunedin.
- H. van Ditmarsch, W. van der Hoek & B. Kooi, 2003, 'Concurrent Dynamic Epistemic Logic', in V.F. Hendricks, K.F. Jorgensen and S.A. Pederson, eds., *Knowledge Contributors*, Kluwer Academic Press, 105 -- 143.
- H. van Ditmarsch, W. van der Hoek & B. Kooi, to appear, *Dynamic Epistemic Logic*, Kluwer-Springer Academic Publishers, Dordrecht.
- H-D Ebbinghaus & J. Flum, 1995, *Finite Model Theory*, Springer, Berlin.
- J. van Eijck, J. Ruan & T. Sadzik, 2004, 'Action Emulation', CWI and ILLC, Amsterdam & Department of Economics, Stanford.

- R. Fagin, J. Halpern, Y. Moses & M. Vardi, 1995, *Reasoning about Knowledge*, MIT Press, Cambridge (Mass.).
- P. Gärdenfors & H. Rott, 1995, 'Belief Revision', in D. M. Gabbay, C. J. Hogger & J. A. Robinson, eds., *Handbook of Logic in Artificial Intelligence and Logic Programming 4*, Oxford University Press, Oxford 1995.
- J. Gerbrandy, 1999, *Bisimulations on Planet Kripke*, dissertation DS-1999-01, Institute for Logic, Language and Computation, University of Amsterdam.
- P. Girard, 2004, 'DDL versus DEL', Department of Philosophy, Stanford University.
- A. Grove, 1988, 'Two Modelings for Theory Change', *Journal of Philosophical Logic* 17, 157–170.
- J. Halpern & M. Vardi, 1989, 'The Complexity of Reasoning about Knowledge and Time', *Journal of Computer and Systems Science* 38:1,195-237.
- P. Harrenstein, 2004, *Logic in Conflict*, Dissertation SIKS 2004–14, Department of computer science, University of Utrecht.
- V. Hendricks, 2002, 'Active Agents', PHILOG Newsletter, Roskilde. In J. van Benthem & R. van Rooy, eds., special issue on Information Theories, *Journal of Logic, Language and Information* 12:4, 469–495.
- W. van der Hoek & J-J. Meijer, 1995, *Epistemic Logic for AI and Computer Science*, Cambridge University Press, Cambridge.
- M. Hollenberg, 1998, *Logic and Bisimulation*, dissertation, Publications Zeno Institute of Philosophy, vol. XIV, University of Utrecht.
- H. Katsuno & A. Mendelzon, 1992, 'On the Difference between Updating a Knowledge Base and Revising it', in P. Gärdenfors, ed., *Belief Revision*, Cambridge University Press, Cambridge, 183–203.
- K. Kelly, 1996, *The Logic of Reliable Inquiry*, Oxford University Press, Oxford.
- B. Kooi, 2003, *Knowledge, Chance, and Change*, Dissertation DS-2003-01, ILLC Amsterdam & Department of Informatics, University Groningen.
- B. Kooi & J. van Benthem, 2004, 'Reduction Axioms for Epistemic Actions', *Proceedings Advances in Modal Logic 2004*, Department of Computer Science, University of Manchester. Report UMCS-04 9-1, Renate Schmidt, Ian Pratt-Hartmann, Mark Reynolds, Heinrich Wansing (eds.), 197–211
- D. Kozen, D. Harel & J. Tiuryn, 2000, *Dynamic Logic*, MIT Press, Cambridge (Mass.).
- F. Liu, 2004, *Diversity of Logical Agents*, Master's Thesis, ILLC Amsterdam.
- Ch. Loedding & Ph. Rohde, 2003, 'Solving the Sabotage Game is PSPACE-hard'.

Technical Report AIB-05-2003, RWTH Aachen, 2003.

- C. Lutz, 2004, 'Expressiveness and Complexity of the Logic of Public Announcements, Informatics Institute', manuscript, Technische Universität, Dresden.
- R. van der Meyden, 1998, 'Common Knowledge and Update in Finite Environments', *Information and Computation* 140:2, 115–157.
- R. van der Meyden, 2005, 'Model Checking the Logic of Knowledge', Tutorial, First Indian Congress on Logic and its Relationship with Other Disciplines, IIT Mumbai .
- J. Miller & L. Moss, 2003, 'The Undecidability of Iterated Modal Relativization', Indiana University Computer Science technical report.
- R. Parikh, 2002, 'Social Software', *Synthese* 132, 187–211.
- R. Parikh & R. Ramanujam, 2003, 'A Knowledge Based Semantics of Messages', CUNY New York & Chennai, India. In J. van Benthem & R. van Rooy, eds., special issue on Information Theories, *Journal of Logic, Language and Information* 12:4, 453–467.
- M. Pauly, 2001, *Logic for Social Software*, dissertation DS-2001-10, Institute for Logic, Language and Computation, University of Amsterdam.
- J. Plaza, 1989, 'Logics of Public Announcements', *Proceedings 4<sup>th</sup> International Symposium on Methodologies for Intelligent Systems*.
- R. Ramanujam, 2005, 'Closing Remarks', First Indian Congress on Logic and its Relationship with Other Disciplines, IIT Mumbai – IMSc Chennai .
- G. Restall, 2000, *An Introduction to Substructural Logics*, Routledge, London.
- M. de Rijke, 1992, *Extending Modal Logic*, Dissertation, ILLC Amsterdam.
- R. van Rooy, 2003, "Quality and Quantity of Information Exchange. In J. van Benthem & R. van Rooy, eds., special issue on Information Theories, *Journal of Logic, Language and Information* 12:4, 423–451.
- H. Rott, 2001, *Change, Choice, and Inference*, Clarendon Press, Oxford.
- J. Ruan, 2004, *Exploring the Update Universe*, Master's Thesis, ILLC Amsterdam.
- M. Ryan & P-Y Schobbens, 1997, 'Counterfactuals and Updates as Inverse Modalities', *Journal of Logic, Language and Information* 6, 123–146.
- T. Sadzik, 2004, 'Epistemic Update as a Dynamical System', Department of economics, Stanford University.
- K. Segerberg, 1994, 'Belief Revision from the Point of View of Doxastic Logic', *Bulletin of the Interest Group in Pure and Applied Logic* 3, 535–553.
- W. Spohn, 1988, 'Ordinal Conditional Functions: A Dynamic Theory of

- Epistemic States', in W. L. Harper et al., eds., *Causation in Decision, Belief Change and Statistics II*, Kluwer, Dordrecht, 105-134.
- R. Stalnaker, 1999, 'Extensive and Strategic Form: Games and Models for Games', *Research in Economics* 53:2, 93-291.
- C. Stirling, 1999, 'Bisimulation, Modal Logic, and Model Checking Games', *Logic Journal of the IGPL* 7:1, 103–124. Special issue on Temporal Logic, edited by A. Montanari & Y. Venema.
- F. Veltman, 1985, *Logics for Conditionals*, Dissertation, Philosophical Institute, University of Amsterdam.
- F. Veltman, 1996, 'Defaults in Update Semantics', *Journal of Philosophical Logic* 25, 221–261.
- G. Venkateash, 2005, 'Temporal Logic with Preferences and Reasoning about Games',  
First Indian Congress on Logic and its Relationship with Other Disciplines, IIT Mumbai 2005 – IIM Bangalore and Sasken.
- A. Yap, 2004, 'Finding out Who is Lying and Cheating in Games',  
Department of Philosophy, Stanford University.
- A. Zanardo, 2002, 'First-order and Second-order Aspects of Branching-time Semantics'. To appear in *Proceedings of HPLMC-02* (Second International Workshop on the History and Philosophy of Logic, Mathematics, and Computation), San Sebastian (Spain), November 7-9, 2002.